NAS2-98005 RTO-77

Technical Research in Advanced Air Transportation Technologies

Single-Year NAS-Wide Benefits Assessment of Multi-Center TMA

Final Report

NASA Technical Monitor:

Daniel Kozarsky

Authors:

Husni Idris Ph.D. - Principal Investigator

Antony Evans

Simon Evans

June 18, 2004

Prepared For:
NASA Ames Research Center
Moffett Field, CA 94035-1000

Prepared By:

Titan Corporation
Air Traffic Systems Division
700 Technology Park Drive
Billerica, MA 01821

Table of Contents

Executive Summary

The objective of this effort is to complete a refined benefits assessment of applying McTMA to meter the arrival flows into the Philadelphia TRACON (PHL) and the New York TRACON (N90). Specifically, the goals are to define and quantify the potential NAS-wide benefits of McTMA for a recent year and year 2015, and to develop a methodology that NASA can use to extrapolate these benefits to other years and other sites for a life-cycle cost-benefit assessment of McTMA. The following report presents the final results of this study.

General Approach

Figure 1 illustrates the general approach by which the benefits of McTMA are identified and analyzed. First, in order to identify the benefits of McTMA the current operations at PHL and N90 are assessed to determine their constraints. The McTMA functionality is also assessed. Then the McTMA benefit mechanisms are identified by applying the McTMA functions to alleviate the identified constraints of the current operations.

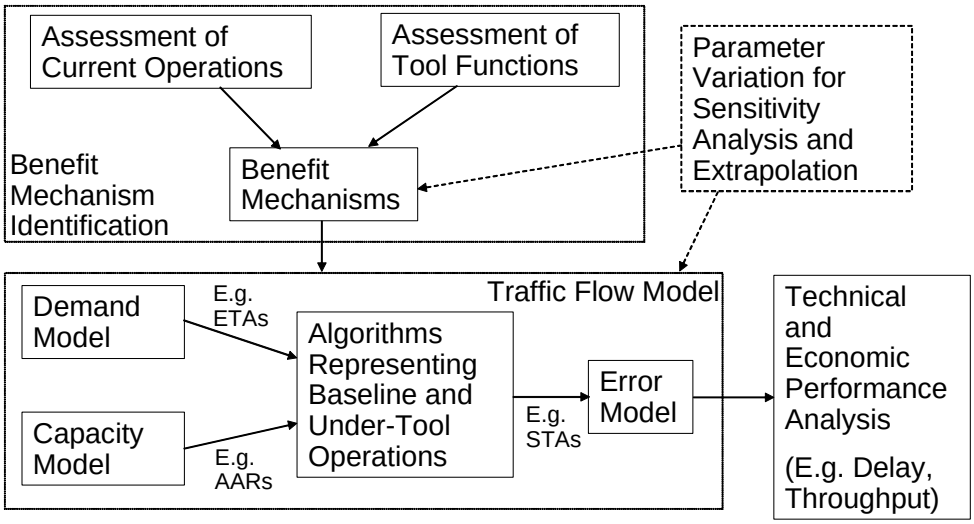

Figure 1. General benefit analysis approach.

After identification of the key benefit mechanisms of McTMA relative to current operations, the current (baseline) system and that as enhanced by McTMA are modeled using a traffic flow model, such that the benefit mechanisms are represented and analyzed. The flow model includes models of the system demand and of the system capacity, which are input to metering algorithms representing the baseline and McTMA operations. The output of the metering algorithms is adjusted through an error model that represents deviations from the desired output due to human and other error sources.

The traffic flow model provides the mathematical abstraction and quantitative metrics in order to measure and compare the performance of the two systems and quantify the benefits of McTMA. An analytical framework based on queuing systems is presented in Section 2. This framework provides a terminology and a quantitative framework to model the baseline constraints, the McTMA functions, and the benefit

mechanisms. Therefore, the current operations and constraints described in Sections 3 and 4, the McTMA functions described in Section 5, and the benefit mechanisms described in Section 6, are modeled quantitatively in the context of this queuing analytical framework, in each section.

As shown in Figure 1, the output of the flow model is then analyzed in order to determine the system technical performance in terms of delay, throughput, and fuel burn. The technical performance of the system is then converted into economic terms, and the economic benefits of McTMA measured. The benefit estimates are finally tested through sensitivity analysis and extrapolated to future years and to other sites.

Identification and modeling of traffic demand, capacity, and traffic management procedures

In order to identify the constraints in the current operations a careful assessment was required of the arrival flows into PHL and N90 and of the current procedures followed in managing and metering these flows. This assessment also led to modeling the arrival demand, modeling the capacity of the system, and modeling the current operations baseline, as described in Section 3.

While the PHL airspace and arrival flows were elaborately studied and described in previous efforts (for example, RTO16 and RTO33 [7]), the N90 airspace and arrival flows were not studied as elaborately prior to this effort. Because it was essential to understand the current flows and operations in order to conduct this study, a significant effort was expended on analyzing the N90 flows based on documentations (such as Standard Operating Procedures of N90 and the surrounding ARTCCs ZNY, ZBW, ZDC, and ZOB and Letters of Agreements between these facilities) and based on information collected through expert elicitation at each of these facilities. Site visits were conducted in the week of November 18th 2002, to each of the above facilities and Traffic Management Coordinators (TMCs) were interviewed.

Demand modeling

These efforts resulted in identifying the arrival flows to PHL and the four major N90 airports analyzed (JFK, EWR, LGA, and TEB) and the meter fixes along these flows where McTMA sequencing and scheduling of aircraft may be applied. One example of the flow networks developed is presented in Figure 2 below. The network includes 3 tiers, each presented in a different color. The arrows in the figure represent the flows modeled. Meter fixes on the flows were specified according to an approach suggested by NASA McTMA researchers. This approach consisted of the specification of meter fixes as close to sector boundaries as possible, but also at flow merge points. Meter fix arcs were also specified in some cases, instead of meter fix points, to ensure that as much traffic as possible is metered, as shown in the figure. The flows are presented for all five airports analyzed in this study in Appendix A, along with further details on how they were generated. These flow networks constituted the underlying structure for modeling the arrival demand.

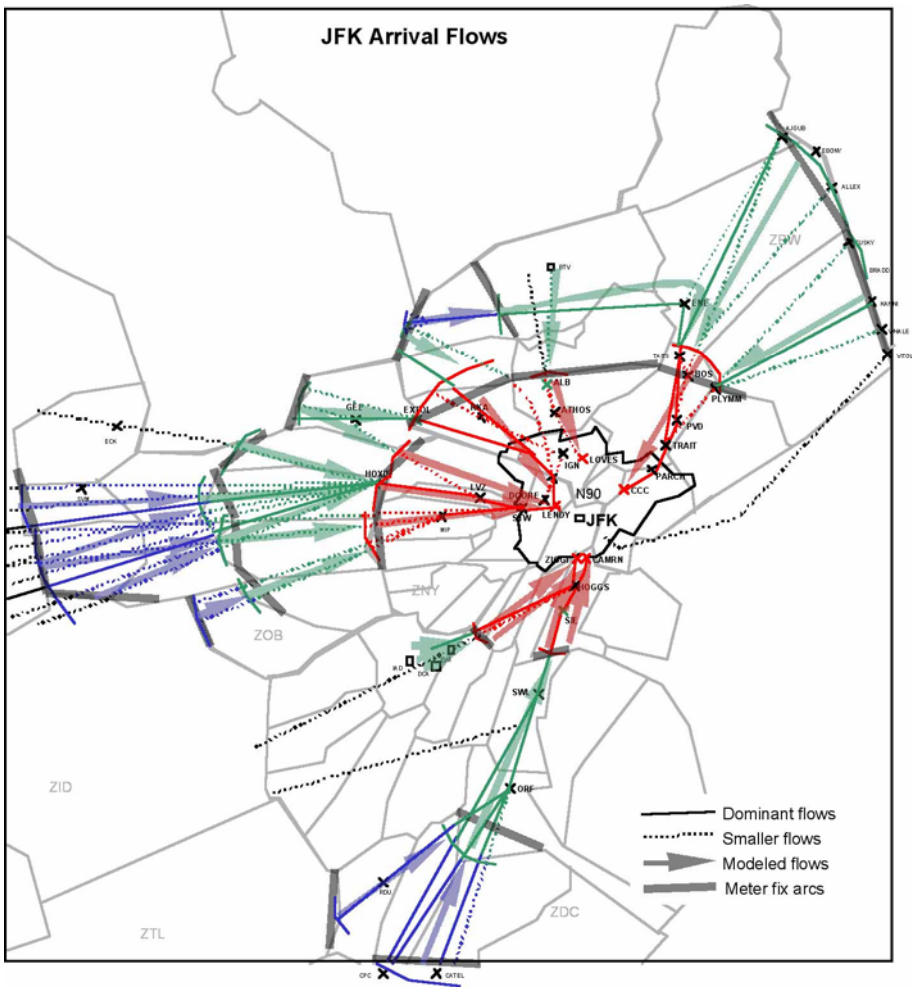

Figure 2. Arrival flow network into JFK

Using historical aircraft position information from ETMS and Host radar tracking data, statistical models of unimpeded transition times between successive meter fixes were generated and used to estimate times of arrival at the meter fixes and at the runways. These statistical models were generated using subsets of flights that encoutered no or small queues and took into account wind, aircraft type, and airport runway configuration. Variability in the estimated times of arrival was modeled by sampling expected unimpeded transition times for each flight from the statistical distributions determined. The flow networks and estimated times of arrival constituted demand input to the algorithms representing McTMA and baseline metering functions and a basis for computing delay in both actual traffic and simulated operations.

These statistical models for unimpeded transition time were tested against trajectory based unimpeded transition times. The trajectories were generated using the CTAS Trajectory Synthesizer process, which is the model used for generating TMA and McTMA estimated times of arrival. These trajectories are generated according to a flight plan and a performance model for each aircraft, and taking wind into account. The statistical models performance compared well to the trajectory based model performance,

with the median of the differences between the models being between 0.0 and 1.6 minutes, thus giving confidence in their representation of the basis for computing delays.

Capacity modeling

The assessment of the current system constraints included a detailed quantitative analysis of the capacities of the 5 airports studied. This analysis was completed using data from the FAA's Aviation System Performance Metrics (ASPM) database.

In order to determine the arrival acceptance rate (AAR) of a runway system, throughput is plotted against demand for each of the commonly reported acceptance rates at each airport as shown in Figure 3. It is clear that the throughput increases linearly with demand, until a maximum is approached, at which point throughput drops off as demand increases. The saturation level is set mainly by the safety separation requirements between aircraft and by controller workload, and is used as an estimate of the maximum service rate capacity of the runway system. The drop may be due to, among other factors, controller workload constraints and airspace complexity constraints. In order to estimate the actual arrival rate capacity of the airport, a hyperbolic curve is fitted to the average throughput with demand less than the drop off point (the dashed line in Figure 3). The second plot in Figure 5 showing the frequency of the demand in 4 aircraft per hour bins shows that the majority of data points fall to the left of this drop-off in throughput. The hyperbolic curve fit asymptotes to throughput equaling demand on the left and to a maximum throughput on the right. This maximum throughput represents the actual capacity operated at the airport, for the reported AAR under question. Because the hyperbolic curve is fitted to the average arrival throughput, this actual capacity represents an arrival capacity average over a range of other varying factors including, for example, departure rate.

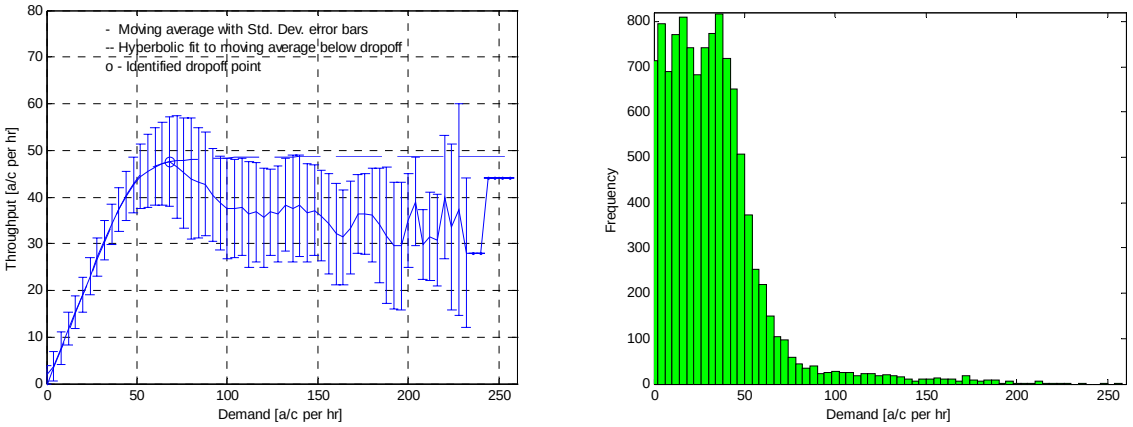

Figure 1. Average throughput versus demand for Newark airport with a reported AAR of 44 aircraft per hour; June, July and August, 2001.

The ASPM capacity analysis also resulted in airport capacity envelopes such as that presented in Figure 4 being identified for the most common configurations operated at each airport. These capacity envelopes plot the arrival throughput versus the departure throughput for each time period and represent the dependence of the arrival rate capacity

on the departure rate. The capacity envelopes also represent other arrival capacity constraints such as the wake vortex separation requirements and controller workload. A number of percentile capacity envelopes are shown in Figure 4. These capacity models were used as an input to the algorithms representing McTMA and baseline metering functions. The 99th percentile measured from half hour throughput data was considered in this study as a conservative maximum safe capacity limit for a runway configuration, as described later and in Section 7 of the report. The asymptotic capacity, computed from the hyperbolic fit described in Figure 3 above and representing the current operated capacity for the runway configuration, is superimposed in Figure 4 showing how it compares to the different capacity envelope percentiles. Details of the capacity analysis are given in Section 3 and Appendix B.

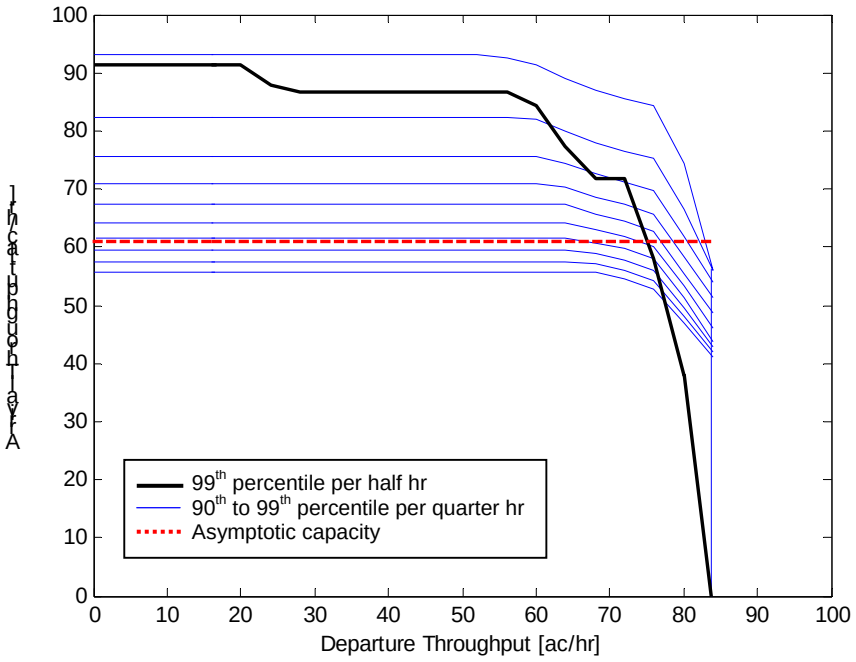

Figure 4. Capacity envelope for PHL with a reported AAR of 52 ac/hr, showing the asymptotic capacity envelope (thick dashed red), modeled capacity envelopes for the 90th to 99th percentiles capacities per quarter hour (thin blue), and the 99th percentile capacity envelope per half hour (thick black).

One observation made in the capacity analysis was that the N90 airspace often constitutes the main flow constraint (bottleneck) to the arrival flows into the N90 airports. This observation was indicated by the N90 TMC during the site visit and was also observed in the capacity analysis presented in Section 3.1. Using historical ASPM data analysis it was observed that the four major airports of N90 (JFK, EWR, LGA, and TEB) are underutilized as a whole (when considered as one landing resource) more than each airport is underutilized separately. However, because McTMA uses airport acceptance rates (AAR) as the applied constraint, an N90 arrival rate capacity was not applied as a constraint to either the baseline or McTMA models in this study.

Baseline modeling

A baseline model that represents the current traffic management procedures was developed and used to measure the benefits of McTMA over current operations in 2003 and in future years. This model consisted of a Miles In Trail (MIT) restriction generation component, a MIT spacing component and a TRACON delay component. Each of these components were derived from and calibrated against actual operation data from facility logs and traffic tracks for November 2003, as detailed in Section 4.

The MIT restriction generation model computed MIT restrictions imposed at the arrival fixes based on airport demand and capacity, and MIT restrictions propagated to upstream boundaries between facilities based on fix demand. The timing and value of the MIT restrictions were derived from analysis of one month of facility logs (November 2003) and were also based on procedures elicited from TMCs during the site visits. They were used to predict the MIT restrictions that would be imposed in future years for the extrapolation of the McTMA benefits.

The MIT spacing model calculated the delay resulting from imposing the MIT restrictions. The MIT restrictions were obtained from the facility logs for the current year and from the MIT restriction generation model for the future years. The MIT spacing model represented deviations from the MIT restrictions imposed as determined from the model calibration against the delay measured in actual operations.

The TRACON delay model calculated the delay imposed in the TRACON in order to satisfy the airport AAR. The traffic enters the TRACON already separated by the MIT restrictions according to the MIT spacing model and any additional delay needed to be absorbed in the TRACON is added by the TRACON delay model. The AAR for each runway configuration was calculated based on calibrating the modeled delay and throughput against those observed in actual operations.

Identification and modeling of McTMA functions and benefit mechanisms

In order to identify the benefit mechanisms of McTMA a careful assessment of the McTMA functionality was accomplished by analyzing McTMA literature and through consultation with NASA's McTMA researchers. The McTMA functions identified are described in Section 5 of this report along with the algorithms that were developed for their simulation. While benefit mechanisms were derived for all the identified functions, only the most important functions (based on NASA's feedback) were modeled and their benefits assessed quantitatively. The functions that were ultimately modeled are:

1. Time based metering, with "delay feedback" and "capacity distribution"
2. Dynamic metering
3. Tiered metering
4. Multiple facility coordination
5. Internal departure scheduling

The benefit mechanisms of McTMA were then derived by applying each of the McTMA functions to alleviate identified constraints and limitations in the current operations. In order to achieve clarity, consistency, and completeness in identifying the benefit mechanisms, formal definitions of functions, benefits, and benefit mechanisms,

and a formal procedure for mapping benefits of McTMA functions were established. The mapping of functions into quantifiable benefits was presented in charts in order to facilitate review by NASA's McTMA researchers. An example chart is given in Figure 5 below. The derivation of benefit mechanisms is described in Section 6 of this report.

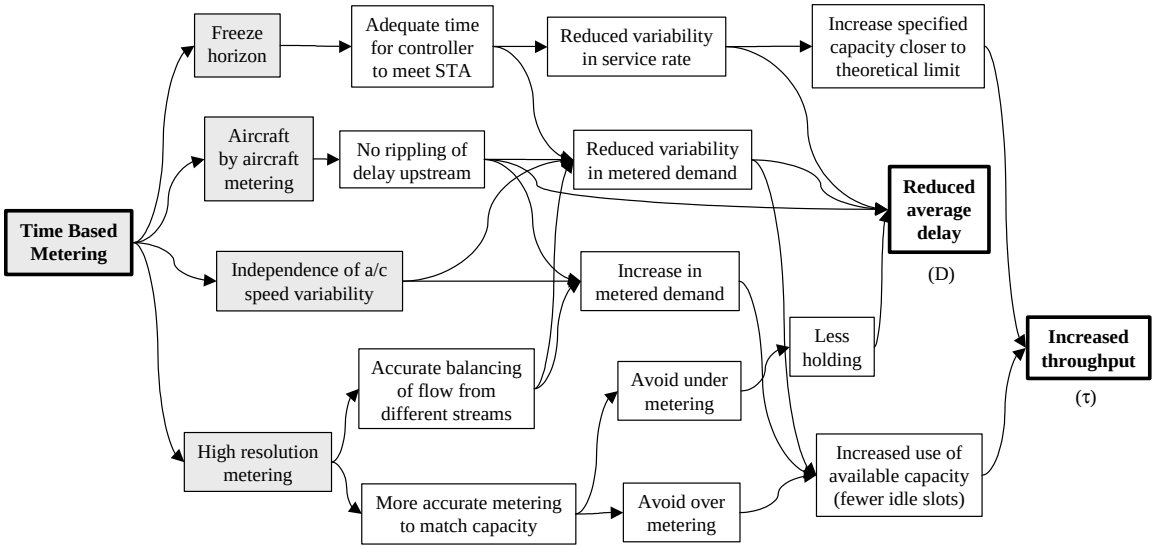

Figure 5. Benefits of time based metering

Conservatively and due to time and resource limitations, not all of the identified benefit mechanisms were modeled and analyzed in this study. The benefit mechanisms not analyzed are indicated with dashed arrows in the benefit mechanisms charts in Section 6. The benefit mechanisms analyzed included mainly:

- Time based metering with its high resolution, no upstream rippling of delay, and accurate balancing between flows from different streams.
- Delay feedback to higher upstream altitudes
- Capacity distribution through distributed scheduling between multiple tires
- Reduced errors in meeting the STAs due to freezing the STAs within a freeze horizon and affording controllers more time and less workload to better meet the STAs
- Higher capacity constraints (closer to the available maximum capacity of the airports) are allowed through a number of the benefit mechanisms as indicated in the charts in Section 6
- Dynamic metering leading to canceling the effect of accumulated errors in meeting the STA between multiple tiers
- Improved balancing between arrival flows through improved coordination between multiple facilities in multiple tiers
- Improved timing of applying metering through better visualization of demand and capacity and better prediction of demand and delays
- Delay of internal departures on the ground before takeoff

The benefit mechanisms not analyzed included most notably:

- Switching flights between arrival flows through improved coordination and offloading between facilities

- Improved decisions to shut off the traffic in extreme situations causing no-notice holding
- Interaction with facilities outside of the McTMA system such as the Command Center or other ARTCCs to reduce the use or severity of ground delay programs or MIT restrictions upstream of the McTMA system
- More optimized delay on the ground for internal departures

An algorithm that mimics the TMA sequencing and scheduling algorithm [4] was used to assign scheduled times of arrival (STA) to aircraft at meter fixes. Flights are assigned scheduled times of arrival at meter fixes in the flow networks to reduce airport arrival throughput below the applied airport capacity and to feedback delay to upstream tiers when the delay absorption capacities of sectors along the route are reached. An error in meeting the scheduled times of arrival was also modeled, according to a normal distribution centered on zero and with a standard deviation of 90 sec. The benefit mechanisms were represented through certain model parameters as described under each benefit mechanism in Section 6.

Technical and economic benefits

The technical performance benefits of McTMA with regard to delay, throughput, and fuel burn savings were estimated relative to baseline delay, throughput, and fuel burn. Section 7 presents these technical performance benefits. In Section 8 of the report the technical performance benefits identified in Section 7 are converted into economic terms.

Fifteen days from November 2003 were simulated, the benefits were calculated per flight and per day, and were extrapolated to the year according to when demand exceeds capacity. This is because benefits are a function of how much metering is applied, on any given day, and flights are metered under McTMA operations when demand exceeds capacity. The fifteen days were selected according to data completeness and represented a random and wide range of metering conditions (demand exceeding capacity).

McTMA generates STAs based on satisfying an AAR constraint, which was one of the key parameters in the McTMA simulation. The arrival rate constraint applied in each 15 minute period, given the number of departures in the period, was read off the capacity envelopes relating arrival and departure rates described in Section 3.1.1. The capacities imposed on the baseline TRACON model (the capacity envelope percentiles that resulted in the best baseline model calibration as described in Section 4) represent the capacity operated under current procedures. McTMA is, however, expected to result in an increase in the throughput by increased utilization of the available maximum safe capacity limit, as described in Section 6.2. A range of benefits were thus calculated by applying a range of capacities above those that calibrated the baseline model, thus modeling benefits due to an increase in capacity utilization under McTMA operations.

It was assumed that McTMA will be limited by runway safety requirements and local operational constraints such as gridlock and controller workload. It was assumed that the 99^{th} percentile of the capacity envelopes calculated per half hour (as described in Section 3.1.1) represent the maximum safe throughput limit for each runway configuration. This maximum is determined by the number of runways, the runway

configuration and the wake vortex separation requirement. The 99th percentile is a conservative choice because the 100th percentile may include possible violations of the safety requirements due to controller human error. It is also conservative because it averages the throughput data available in 15 minute intervals, reducing the binning error. It was assumed that this maximum safety limit will not be increased due to the application of McTMA's time based metering since McTMA will not impact the number of runways and wake vortex separations. Any applied capacity constraint for which this maximum safety constraint was not violated by the arrival throughput any more than in actual operations (calculated from actual landing times) was considered to meet the safety requirements. The throughput was permitted to violate the safety constraint to the same extent as actual operations allowing the same level of safety as currently practiced. The highest percentiles of the capacity envelopes that met this criterion were applied and the corresponding benefits considered the highest benefits achievable by McTMA. This high limit does not necessarily represent what controllers may achieve in practice under McTMA operation, but rather only what is possible and available to achieve.

Figure 6 displays the benefit estimates at each airport for the range of capacity constraints applied (as percentages of the capacity envelopes at each airport) between the capacity that calibrated the baseline model (representing no capacity increase due to McTMA) and the 98th percentile of the capacity envelopes. The nominal benefits case corresponding to the maximum safe capacity limit that McTMA may impose at each airport is indicated on each range with a circle. Numbers are shown in Table 1 corresponding to the minimum (no capacity increase) and nominal yearly expected benefits under McTMA.

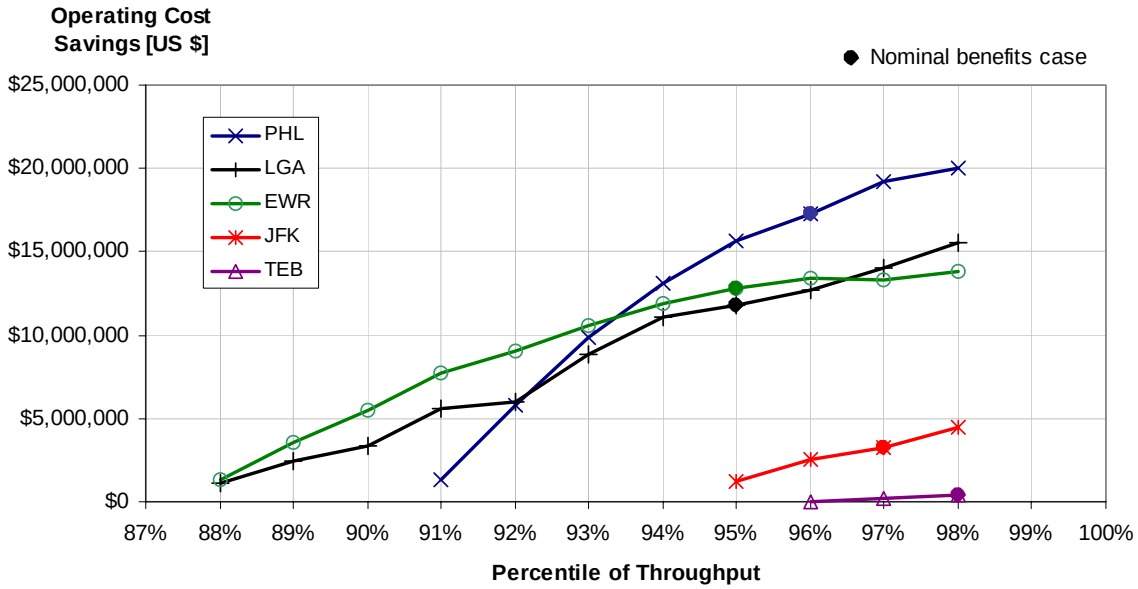

Figure 6. Yearly increase in operating profit due to McTMA, 2003 – varying applied airport capacity under McTMA.

Table 1. Yearly increase in operating profit due to McTMA, 2003.

Airport	2003 Yearly Savings [US$ / year]	
	No Capacity Increase	With Capacity Increase
PHL	1,302,000	17,275,000
LGA	1,141,000	11,810,000
EWR	1,343,000	12,817,000
JFK	1,268,000	3,289,000
TEB	4,000	383,000

Another key parameter in estimating the McTMA benefits was the duration of metering periods over which time based metering was applied in the McTMA simulation and MIT was applied in the baseline simulation. According to consultation with McTMA researchers time based metering was applied when demand was expected to exceed the reported capacity of an airport. However, as expected to be the case in McTMA operation, delay savings were counted only after the delay requirement exceeded two minutes (which may be absorbed within the TRACON) and until the delay requirement is reduced to zero. Minimum metering durations of 30 minute and minimum time separation between metering periods of 30 minute, were imposed. In the baseline model MIT were applied during the MIT periods according to the facility logs.

Sensitivity analysis

The sensitivity of the results in Table 1 to changes in a number of parameters relating to the benefit mechanisms and to the modeling was studied. These parameters were arrival capacity (shown in Figure 6), sector capacities (for feedback of delay upstream under McTMA operations), error in meeting scheduled times of arrival, operating cost per flight, and metering periods, including the minimum duration of metering, and the criteria for initiating metering. The results showed high sensitivity to the applied arrival capacity, the mean error in meeting scheduled times of arrival, and the criteria for initiating metering. The results of the sensitivity analysis are presented in Section 9.

Extrapolation to future years and other facilities

Extrapolation to benefits for 2010, 2015 and 2025 were completed according to forecast demand in these future years. Baseline operations in the future years were modeled by applying increased MIT restrictions, as calculated from the MIT generation model, to the increased arrival demand. The demand was increased maintaining the temporal and spatial dynamics of the current schedule at each airport. The results are presented in Table 2. The results for LGA plateau because forecast demand is capped after 2005.

Table 2. Yearly increase in operating profit due to McTMA, extrapolated to 2010, 2015, 2025 (in 2003 US Dollars).

Airport	Yearly Savings [2003 US$ / year]			
	2003	**2010**	**2015**	**2025**
PHL	17,275,000	56,538,000	94,779,000	363,808,000
EWR	11,810,000	28,569,000	29,480,000	29,409,000
LGA	12,817,000	96,655,000	196,430,000	520,122,000
JFK	3,289,000	14,595,000	40,253,000	172,494,000
TEB	383,000	605,000	1,635,000	12,922,000

It is important to note that the delay levels for the simulated 2015 results were very high under both current operations and McTMA operations. This is because the demand increase forecast by the FAA APO TAF [17] is not feasible given the capacity applied, under current operations or McTMA operations. This observation is consistent with [18], which suggests that current demand forecasts do not adequately account for capacity constraints.

A model for extrapolating the benefit estimates to other McTMA sites was also built as shown in the equation below. The model is based on queuing dynamics relating the benefits (or delay savings) to the number of delay operations (*ArrOps*), the percentage of time demand exceeds capacity ($P_{D>C}$) and the utilization of available capacity (*Util*) at an airport. The model parameters were selected such that they may be calculated from available data sources such as ASPM, and was calibrated against the five airports analyzed to determine the constant values.

$$Benefits = 19{,}366 \times \left[\frac{ArrOps_{perday}}{1-(0.0213)(P_{D>C})(Util)} \right]$$

Final remarks

The benefit estimates of McTMA assessed in this study are believed to be realistic, robust, and conservative for a number of reasons:

1. A large sample of days was analyzed representing a random and wide range of metering conditions. Days or periods of time when the system was thought to be restricted by constraints irrelevant to McTMA (or not included in this study) such as local restrictions not related to runway capacity or strategic restrictions like ground delay programs, were excluded. This was done in order to limit the benefits assessment conservatively to those delays that McTMA is believed to mitigate.
2. The benefit estimates resulted from comparing a McTMA model of time based metering to a model of baseline operations using distance based metering. Care was

taken to model in the baseline only procedures and dynamics that are relevant to McTMA (such as MIT delay propagated from the runway and not local MIT, local rerouting, or GDP).

3. The delays were measured with respect to unimpeded estimated times of arrival derived from statistical models based on analysis of historical track data. The statistical unimpeded times took wind, aircraft type, and runway configuration into account. They compared well to estimated times of arrival computed from trajectory synthesis (the CTAS Trajectory Synthesizer process) based on flight plan, wind, and aircraft performance.
4. The benefits assessment focused on a subset of McTMA functions as described in Section 5.
5. Due to time and resource limitations, not all of the benefit mechanisms were modeled and analyzed in this study. The benefit mechanisms not analyzed are indicated with dashed arrows in the benefit mechanisms charts in Section 6.
6. Care was taken to make conservative assumptions about the McTMA operation in the field and about modeling parameters, and to consult NASA's McTMA researchers and their experience with the tool and in the field. Through sensitivity analysis (described in Section 9) a range was tested for many of these assumptions and parameters to provide a range of corresponding benefit estimates and to assess how much of an impact such assumptions and parameters have on the benefit estimates.

Acknowledgments

The Titan team would like to extend their appreciation to a number of groups and individuals of the NASA Ames Research Center and the FAA for their help and assistance in this work. Namely:

Thanks to Michelle Eshow, Akbar Sultan, and the CTAS software group for providing the Titan team with a suite of MATLAB tools that processes and analyzes Host track data. The MATLAB suite proved extremely useful. The Titan team built upon this MATLAB suite and added a number of analysis tools that will be provided to NASA for their future use.

Thanks to Ron Reisman, Matt Ma, Todd Farley and any other NASA researchers who helped in providing Host, ETMS and RUC data.

Thanks to Monicarol Nickelson for her assistance in obtaining Command Center and other facility logs from the FAA. And thanks to the FAA for collecting and providing the logs.

Thanks to William Chan for his assistance in calculating and providing fuel burn data.

Thanks to the NASA and FAA McTMA teams for organizing and participating in site visits to N90, ZNY, ZBW, ZDC, and ZOB. Special thanks go to Todd Farley for leading this effort.

Thanks to the FAA traffic mangers of N90, ZNY, ZBW, ZDC, and ZOB, for providing invaluable information, standard operating procedures, and letters of agreements, during the site visits, and for collecting and providing facility logs for the month of November 2003.

Thanks to Dave Knorr and the FAA free flight office members who attended two briefing of this work and provided invaluable feedback.

Thanks to the NASA McTMA and TMA researchers for providing input, feedback, and support throughout this work. Special thanks to Tom Davis, Todd Farley, Steve Landry, Cheryl Quinn, Katharine Lee, Ty Hoang, Gregory Wong, Steve Green, and Parimal Kopardekar. Special thanks to Todd Farley and Steve Landry for their continuous availability to provide feedback throughout the project.

Finally, special thanks to Dan Kozarsky, NASA's technical monitor of this task, for his invaluable support, encouragement, direction, and feedback.

The Titan team would also like to thank their Titan colleagues who supported this project at different times, including: David Chesler for his support in software work related to the CTAS Trajectory Synthesizer process, Ted Hsu for his help in analyzing MIT restrictions, to Bob Vivona for his support in providing domain expertise related to TMA and McTMA, and to Allan Krueger for his support as a program manager.

This work was performed under contract number NAS2-98005 RTO-77.

1. Introduction

The Traffic Management Advisor (TMA) is a decision support tool – one of the Center TRACON Automation System (CTAS) suite of decision support tools. It is developed by NASA Ames' Advanced Air Transportation Technologies (AATT) Project to improve the performance of arrival flows into congested terminal areas. The underlying concept of TMA is a time-based metering technique that generates a sequence, a schedule, and a runway allocation for aircraft arriving at TRACON feeding gates [1,2,3,4]. TMA has successfully been implemented for a number of Terminal Radar Controls (TRACONs) fed by a single Air Route Traffic Control Center (ARTCC) [1,3].

Multi-Center TMA (McTMA) is an extension of TMA designed for a TRACON with multiple feeding centers, which is particularly characteristic of the US northeast airspace [2,3]. McTMA is essentially a network of communicating TMAs that are deployed at the multiple ATC facilities that feed a particular TRACON and are modified to accommodate the multiple-center environment. McTMA is currently being implemented at Philadelphia TRACON (PHL), which receives traffic from New York ARTCC (ZNY) and Washington ARTCC (ZDC) [2,3]. McTMA implementation is expected at New York TRACON (N90), which receives traffic from ZNY, ZDC, and Boston ARTCC (ZBW).

In addition to dealing with metering TRACON arrival flows being fed from multiple centers, McTMA extends the metering horizon upstream to multiple centers when needed. This is required in the case of the PHL and N90 (northeast) environment, where some centers feeding the TRACON, particularly ZNY, are small in size and do not provide enough time and space horizon to meter the arrival flows into PHL and N90 [5]. McTMA in this case starts metering the PHL and N90 arrival flows from the time they enter ZOB or from within ZOB. Therefore, the McTMA environment may involve multiple centers extending multiple tiers from the destination airport or TRACON, as well as multiple centers within each tier.

2. Approach and Methodology

Figure 1 illustrates the general approach by which the benefits of McTMA are identified and analyzed. First the McTMA benefit mechanisms are identified as described in Section 2.1, then the performance of the operations under McTMA and under current operations are modeled using a traffic flow model and compared, as described in Section 2.2. Finally the technical performance is converted to economic terms, sensitivity analysis and extrapolation of the benefits to future years and other McTMA sites are performed as described in Section 2.3.

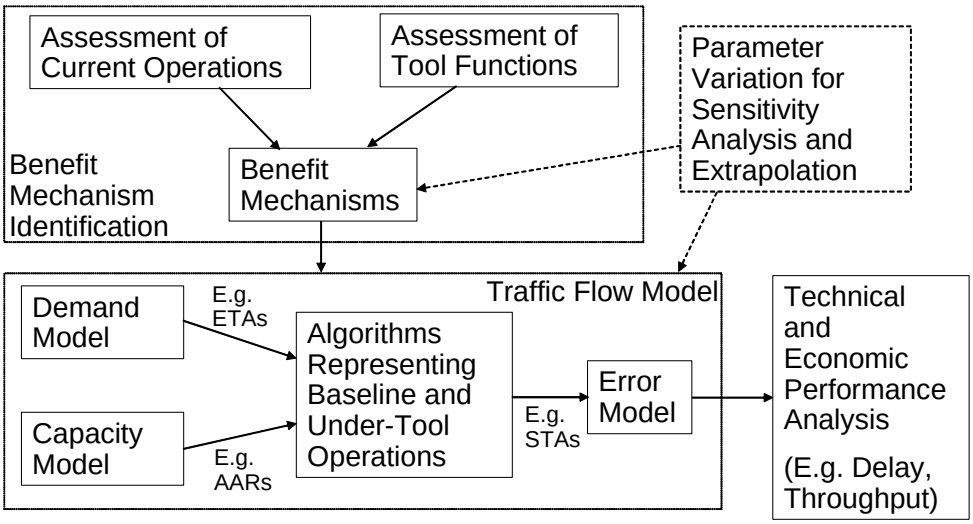

Figure 1. General benefit analysis approach.

2.1. Identification of Benefit Mechanisms

First, the benefit mechanisms of McTMA are identified. In order to ensure that as many benefits are captured as possible, the applicability of the benefit mechanisms identified, and feedback from McTMA researchers, it is essential that a formal, reviewable approach to the identification of the benefit mechanisms be developed. Therefore, for the purpose of clarity, consistency, and completeness, functions, constraints and benefits are formally defined as follows:

- A *Function* is a user utility of the tool.
- A system *Constraint* is any condition that causes demand to exceed capacity of a NAS resource.
- A *Benefit* is a quantifiable performance advantage or operational enhancement that has a direct stakeholder impact.
- An *Economic Benefit* is a benefit directly quantifiable in monetary terms, and leads directly from a *Benefit*.

- A *Benefit Mechanism* is a linkage that converts a function into a benefit by applying the *function* to alleviate system *Constraints*.

A function excites a benefit mechanism, which creates a benefit.

Based on these definitions, the benefit mechanism identification approach includes the following primary components:

1. Identification of current operations, including identification of current system flow constraints and flow management procedures
2. Identification of the McTMA functionality
3. Identification of the benefits of each McTMA function by applying the function to alleviate the identified system constraints and limitations of the current operations. This includes mapping separate benefit mechanisms for each function according to the constraints of current operations.

By identifying McTMA functionality before the identification of benefit mechanisms it is ensured that benefits from all McTMA functions are accounted for. The approach is similar to that used in TO10 [6] and TO33 [7], although it includes a more rigorous identification of benefits. The current constraints, McTMA functionality and McTMA benefit mechanisms identified are detailed in Sections 3, 5 and 6, respectively, along with details of the modeling of each.

2.2. Traffic Flow Model and Baseline Comparison

After identification of the key benefit mechanisms of McTMA relative to current operations, the current (baseline) system and that as enhanced by McTMA are analyzed accordingly, using a traffic flow model. The traffic flow model provides the mathematical abstraction and quantitative metrics in order to measure and compare the performance of the two systems.

The traffic flow model includes algorithms that represent the metering processes of both McTMA and of the current operations. Inputs to the flow model come from a model of the system demand and of the system capacity. The output of the metering model is adjusted through an error model that represents deviations from the desired output due to human and other error sources. The output of the flow model is then analyzed in order to determine the system performance.

2.2.1. Queuing Network Abstraction

In order to provide a quantitative framework with which to analyze and compare the dynamics and performance of the arrival process under McTMA, and under the current system, the arrival process is abstracted to a queuing network. This queuing network forms the basis of the flow model of the arrival process. The network includes multiple tiers of fixes at which metering takes place, feeding into one or more airports. Figure 2 illustrates this schematically for the flows into JFK, through hypothetical meter fixes.

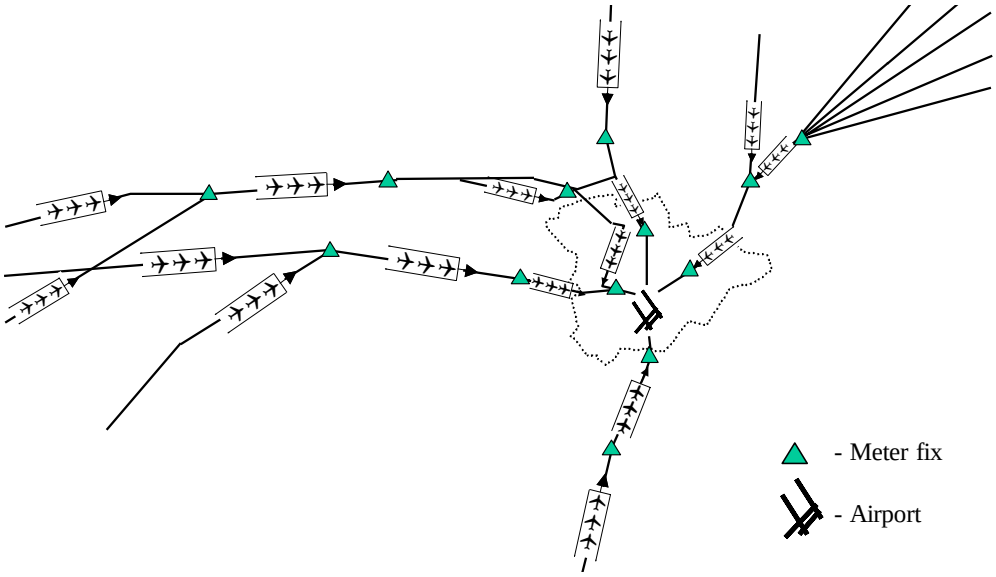

Figure 2. Schematic arrival queuing network for JFK.

The queues represent delay absorbed through ATC interventions such as holding, speed reduction and vectoring. Queues are limited in capacity according to how much delay is able to be absorbed in the available airspace. Metering the flow into a resource by limiting the flow through the upstream resource is used to reduce the adverse effects of high demand or reduced capacity.

In order to illustrate the process of metering the arrival flow and to define key variables that will be used in the system modeling and benefit analysis, a core element of the queuing network is presented in Figure 3. This core element contains a resource that constrains the flow and a resource used to meter the flow, along with their associated queues.

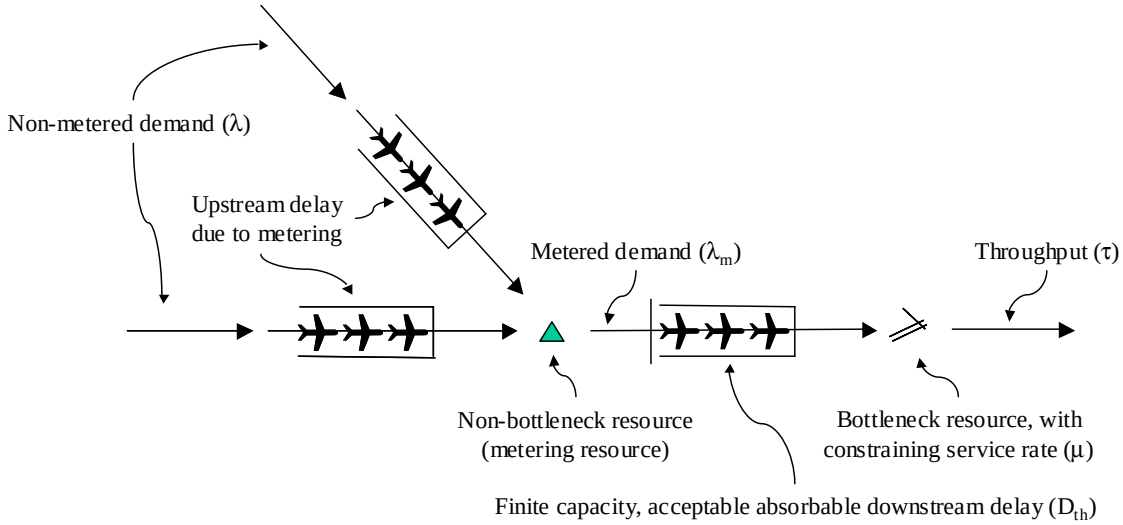

Figure 3. Schematic of queuing network applied to arrival process sub-component.

The system can be described in terms of queuing abstraction as follows:

- A resource, such as a runway or meter fix, can service aircraft at an average service rate μ. This service rate represents the finite capacity of the resource to serve aircraft, due mainly to safety separation and workload constraints. Each resource in the system has a service rate, but under high system demand or reduced capacity, the slowest resource becomes the binding constraint on the flow. This resource is the bottleneck or constraining resource.
- Airline scheduled traffic arrives at each fix, i, at an average arrival rate λ_i. The demand at a fix is typically represented by a series of Estimated Times of Arrival (ETA) calculated based on the airline schedules and flight plans. When an aircraft arrives at a fix at a time later than its ETA it incurs a delay D. Delays occur when the demand is higher than the service capacity.
- Even when the average arrival rate is less than the average service rate, the aircraft may incur a delay due to the variability in the demand (represented by the standard deviation of the inter-arrival times, σ_x) and the variability in the service rate (represented by the standard deviation in the inter-service times, σ_s).
- Delay can be absorbed in the sector airspace between fixes through vectoring, speed reduction and holding. The amount of delay that can be absorbed in any resource is limited to a finite capacity due to workload and safety separation constraints. The capacity of a sector is typically measured by an Operational Acceptable Level of Traffic (OALT)[1]. This number specifies the maximum number of aircraft that can be within the sector at a given time, and corresponds to an acceptable level of absorbable delay. When more delay must be absorbed than this limit, the excess delay needed to be absorbed is propagated to the upstream resource, to be absorbed there. This is known as blocking, in which a resource blocks further acceptance of traffic, leading to the propagation of delays upstream.
- During rush periods scheduled demand is often higher than capacity, due to airline and passenger scheduling preferences. On average, however, demand must be maintained below capacity, so as to avoid an unstable growth in delays. The degree to which the demand matches the capacity can be measured by the ratio of average demand λ to average capacity μ. This is called the utilization ρ of the resource given in Equation (1):

$$\rho = \frac{\lambda}{\mu} \qquad (1)$$

[1] According to correspondence with a TMC, the OALT is a number agreed upon by a member from management (usually the Area Operations Manager) and Area NATCA representatives. The number they come up with is then incorporated into the Monitor Alert; another representation of the OALT is therefore, the Monitor Alert Parameter (MAP). The Monitor Alert Parameters (MAP) can be adjusted by ± 3 minutes by traffic management for various traffic management issues. Many times the numbers identified for a specific sector are incorrect - a known shortfall of the Monitor Alert Program.

- When demand is higher than capacity, metering is a process that attempts to match the demand with the service rate capacity of a downstream resource in order to distribute the delays upstream. The flow is metered from the un-metered demand λ, to a metered demand λ_m. Downstream of the metering resource the arrival rate into the downstream resource is the metered demand λ_m. The metering of the flow at the metering resource is according to constraints fed back from downstream.
- The actual flow through the system, or system throughput τ, is dependent on the demand and metering in the system. When capacity exceeds demand, throughput equals demand. However, when demand exceeds capacity, throughput is instead determined by the capacity of the constraining resource.

If the metering were excessive, the metered demand would also be less than the capacity of the constraining resource μ, leading to lost capacity. However, in order to minimize this lost capacity, the flow can be metered to a level greater than the capacity of the constraining resource μ to maintain pressure on it.

The impact of changes in demand, capacity, their variability and their ratio (utilization), on delay, can be understood by considering the heavy traffic approximation for a G/G/1 queuing system (general inter-arrival distribution/general service distribution/single server), as follows:

$$\overline{W_q} \approx \frac{\lambda_m\left(\sigma_x^2 + \sigma_s^2\right)}{2\left(1-\rho_m\right)} \qquad (2)$$

Where $\overline{W_q}$ represents the average wait time in the queue, or average aircraft delay D. Equation (2) can be used to gain insight into the behavior of the system in terms of tradeoffs between delay and throughput, and to estimate general trends and verify results approximately. Clearly, as metered demand λ_m, or variabilities σ_x or σ_s decrease, the average delay also decreases, as expected. Similarly, an increase in capacity μ (reduction in ρ) results in a decrease in average delay. However, if variabilities σ_x or σ_s decrease, or capacity increases, and the average delay was maintained at the same (acceptable) level, it is possible to increase the demand and hence the throughput of the system.

2.2.2. McTMA and Baseline Comparison Methodology

The queuing abstraction described in Section 2.2.1 provides a common framework to describe and quantify the dynamics and performance of the arrival process, and metering in particular, under both the current operations and under McTMA. When demand is higher than capacity the current ATC system utilizes mainly Miles In Trail, which is a distance-based metering technique that requires aircraft to be spaced, at a specified fix, by a certain number of miles. Alternatively, McTMA provided a time-based metering technique used to sequence and schedule aircraft at meter fixes. The performance of the operations with the use of McTMA's time based metering technique is to be compared to the performance of the current baseline operations using the distance based Miles In Trail.

In order to measure the performance of the operations with the use of McTMA, these operations need to be simulated. An algorithm that mimics the TMA sequencing and scheduling algorithm [4] was used to assign scheduled times of arrival to aircraft at meter fixes. A selection of meter fixes for N90 is suggested in this report (Appendix A) and was adjusted based on NASA's feedback. [2]

The approach in this study was to compare the performance of the simulated operations with the use of McTMA to the performance of the modeled operations under current procedures as a baseline. The baselines that were compared in the study are, therefore, shown in Figure 4. The baseline model of the current operations was generated using actual traffic data for particular days when Miles In Trail were used to meter the arrival flow into PHL and N90. Traffic data (aircraft radar-tracked positions) were obtained from NASA in two forms, Host data and Enhanced Traffic Management System (ETMS) data, for 11 days in August and September of 2002 and for the month of November 2003. Command Center and local facility logs for these days were also made available through NASA in electronic or paper form.

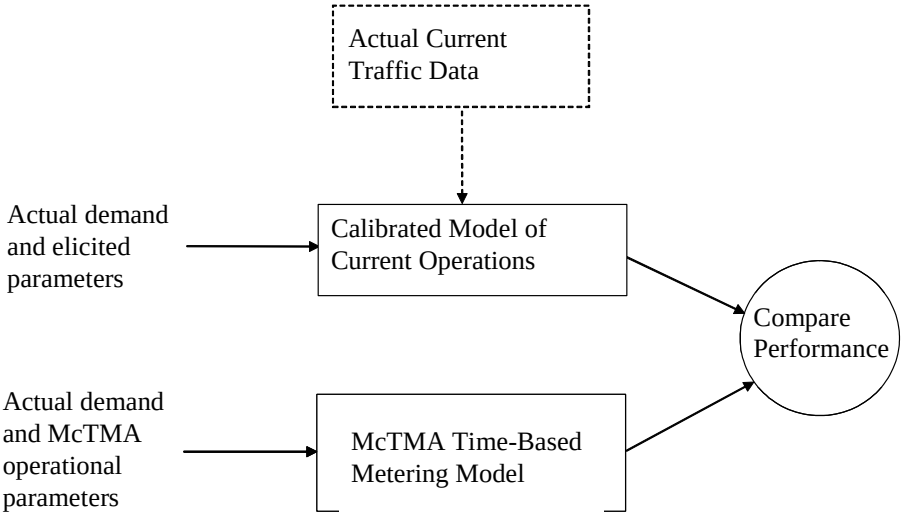

Figure 4. Comparison of McTMA and baseline operations

A baseline model was used rather than actual traffic data because actual data may provide a misleading baseline for a number of reasons. One reason is that the actual

[2] It was intended to analyze two modes of McTMA operation representing an incremental use of the tool. In the first mode McTMA would be used as a demand visualization/situation awareness tool that would be available to TMCs before the ATC site is converted to the use of time-based metering. In this mode the tool would help the TMCs in the selection of more optimal MIT restrictions given the demand predicted more accurately by McTMA. In the second mode the time based metering function of McTMA is used. The first mode of operation would be modeled as a what-if type capability to assist in the selection of MIT restrictions. However, due to time constraints, and because such a function is not currently intended for McTMA [2], neither incrementally nor on a permanent basis, the analysis of this mode of operation was dropped in favor of higher priority analyses.

traffic data represents a manifestation of all the sources of inefficiency in the current operations, not all of which may be addressed by McTMA. If using actual data, care must be taken to exclude the non-relevant sources of inefficiency when selecting actual traffic samples for measuring the baseline performance. Failure to do so renders the estimated benefits optimistic. Another reason that actual traffic data may provide a misleading baseline is the discrepancy between the reported information and the actual performance of the system. For example, the actual data often does not accurately reflect the effect of the specified restrictions reportedly imposed at a given time. In many cases the system outperforms, in terms of delay and throughput, the performance expected given the imposed restrictions. (For example, when 15 miles in trail are specified often 13 miles in trail may be imposed in practice, resulting in a higher throughput than expected have 15 miles in trail been imposed.)[3] Failure to account for such an inaccuracy renders the estimated benefits conservative and often negative. Another difficulty with the actual traffic data is its incompleteness and its inaccuracy in measuring the performance of the system. For example, actual aircraft positions from radar tracking systems may be available through Host computers with12-second intervals, or from ETMS with one-minute intervals. Therefore, Host track data is more accurate in measuring fix crossing times and landing times. Therefore, using actual traffic data as a baseline can result in conservative or optimistic benefits estimates if not selected carefully. For this reason, the current operations were simulated instead. Examples of such an approach include the benefit study of Regional Metering, a potential enhancement to TMA/McTMA, TO71 [8], which used simulated current operations as a baseline instead of using the actual data. TO71 used for the baseline an algorithm that selects optimal MIT restrictions.

Simulation allows concentrating on the elements that are believed to be relevant to the benefits assessment by excluding the sources of inefficiency that may not be mitigated by McTMA. The performance of the actual operations is important for the purpose of calibration of the simulated baseline and for the identification of the available pool of inefficiencies in the current operations and the fraction of it that McTMA mitigates. In order to exclude sources of inefficiency that McTMA may not address, particular days when only Miles In Trail were in effect were selected for calibration. These days (15 in total) represented a variety of demand levels and of constraints severity and included occurrences of holding and thunderstorms.[4] Facility logs indicated what restrictions were imposed on each day. The traffic demand was then run through simulations that represented the operations under current operations, and the performance compared to the actual traffic data for calibration.

2.2.3. Site Visits for Identification of Current Operations

In order to identify the benefits of applying McTMA to the PHL and N90 arrival flows, the current operations and current flow constraints must be well understood and

---

[3] The deviation from the restrictions is due to a number of reasons including human error and the non-dynamic nature of the restrictions.

[4] It was intended to analyze days when other restrictions, such as a Ground Delay Program, were in effect, and to analyze the impact of McTMA on reducing the delays imposed by such more severe flow management restrictions. However, due to time limitations this was not accomplished.

analyzed. As described in Section 2.2.2, the current performance of the system provides a basis for comparison with the system performance under McTMA.

Through site visits to the main ATC facilities that control the flows into PHL and N90 (N90, ZNY, ZBW, ZOB and ZDC[5]) a number of insights were gained into the flows, capacity constraints, and flow management procedures at these facilities. Husni Idris and Antony Evans from Titan, and two McTMA researchers, Todd Farley from NASA Ames, and Steve Landry from Raytheon, participated in the site visits. One facility was visited on each day of the week of November 18th 2002, in the order listed above. At each facility one or two traffic managers were assigned to the team to describe the flows and operations and to address questions. Copies of the Standard Operating Procedures (SOPs) of each facility and the Letters Of Agreement (LOAs) between facilities were obtained. These documents were needed in order to determine the arrival flows into N90 (The PHL flows were determined in previous studies). The arrival flows into PHL and N90 are described in demand modeling (Section 3.2) and Appendix A of this report along with a suggestion for the meter fixes to be used by McTMA for the N90 flows.

The visits were successful in gaining insight into the main flow constraints and the current flow management procedures at each facility. A list of questions and parameters that needed to be identified to support the modeling efforts was prepared beforehand. In some cases it was possible to obtain the experts' estimation of certain numerical parameters such as delay thresholds and rule of thumb procedures for Miles In Trail propagation. However, as expected due to the limited time and resources, the set of modeling parameters and procedures identified was incomplete and was pursued through further elicitations. Contacts were made at each facility in order to ask further questions that were identified at a later time as the modeling and analysis efforts continued. Follow-on questions were sent through NASA's McTMA researchers during their additional visits to the sites. It was possible to get more information from ZNY, but unfortunately, not from ZBW, ZDC, and N90.

While some information were needed for modeling McTMA function parameters, the majority of the information requested was needed for modeling a baseline that represents the current operations. This baseline was used as a current year baseline as well as a baseline for future year benefits assessment. The current operations model was designed to mimic the current behavior of the traffic managers and controllers. Due to the incompleteness of the elicited information the model of current operations was instead based mainly on observations of actual data from November 2003 (track data and facility logs). Data analysis were performed to attempt to identify the way the traffic managers set the system capacities and Miles In Trail restrictions and the factors that they are based on, the way they propagate the Miles In Trail restrictions through the system, and the way they implement (and possibly deviate from) the specified Miles In Trail. These elicited and modeled baseline flow management procedures and behavior are described in Section 4.

[5] PHL was not visited due to ongoing equipment upgrade activities at the facility and the prior knowledge about the PHL operations from previous studies.

2.3. Economic Performance, Sensitivity Analysis, and Extrapolation

As shown in Figure 1, the technical performance of the system is converted into economic terms, and the economic benefits of McTMA measured. The analysis is then tested through sensitivity analysis to identify the sensitivity of the benefit estimates to certain model and benefit mechanism parameters. The analysis is also extrapolated for extension to other years and to other McTMA sites, so that the economic benefits of McTMA can be identified across the NAS, for an extended period of time.

For the assessment of the McTMA benefits in later years, a simulated baseline that represents the N90 and PHL operations without using McTMA in later years was compared with simulated operations with using McTMA in later years. The simulated later years operations accounted for differences from current operations, as possible. These differences were assessed based on FAA studies and documentations. They included increase in demand according to the FAA forecasts and changes in the airspace structure (namely, accounting for the consolidation between N90 and PHL). The addition of runways was also investigated. The extrapolation is described in Section 10.

3. Identification and Modeling of Capacity and Demand

The capacity constraints and demand flows for PHL and N90 are presented in Sections 3.1 and 3.2, respectively, along with the corresponding models.

3.1. Capacity Constraints

Aircraft transition from en-route sector airspace to terminal area airspace (TRACON), and then land on specific runways at the destination airport. The capacity constraints of these three major resources (runway, TRACON airspace, and sector airspace) are described below in Sections 3.1.1 to 3.1.3.

3.1.1. Runway Capacity Constraints

The primary flow constraint is usually the airport acceptance rate, which depends mainly on the runway configuration, visibility, and runway conditions. The acceptance rate of an airport is usually reported by the airport control tower and changes depending on changes of runway configuration and airport conditions. These reported acceptance rates however, are inaccurate as they represent a crude and conservative estimate, and the actual operations on any particular day may deviate largely from them.

In order to determine the arrival service rate (known as the Arrival Acceptance Rate, AAR) of a runway (or runway system), the throughput of the runway system is plotted as a function of the demand, as shown in Figure 5. Because the arrival service rate is a function of runway configuration, visibility, and runway conditions, and because the reported acceptance rate is generally adjusted according to each of these constraints, throughput is plotted against demand for each of the commonly reported acceptance rates at each airport.

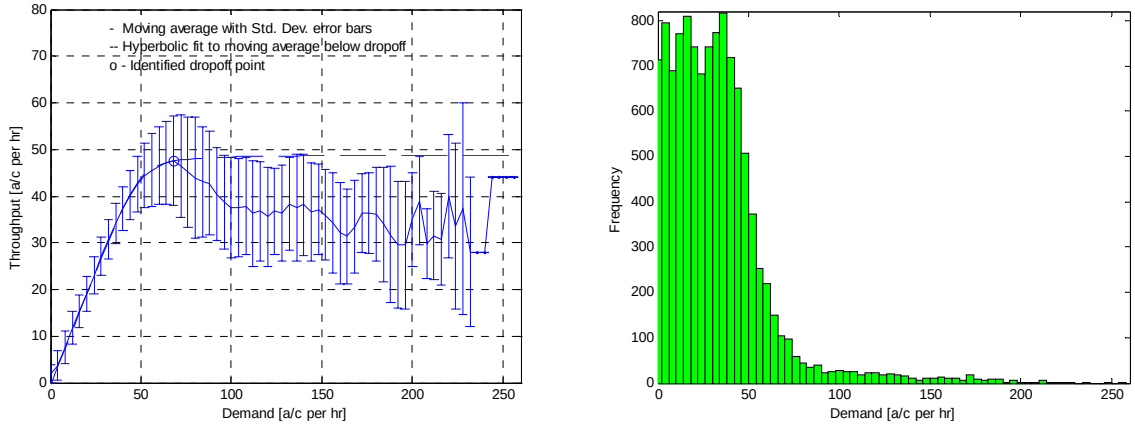

Figure 5. Average throughput versus demand for Newark airport with a reported AAR of 44 aircraft per hour; June, July and August, 2001.

As shown and described in detail in Appendix B, such charts were generated using throughput and demand data for each quarter hour period from June to August of 2000, and January to August of 2001. This data is from the FAA's Aviation System Performance Metrics (ASPM) database. Throughput represents the number of aircraft that landed at the airport per quarter hour, multiplied by 4 to be specified as an arrival rate per hour. Demand, as recorded in the ASPM dataset, is the number of aircraft that intend to land at a specific airport in a unit of time. An aircraft is included in the demand in the periods starting from its actual wheels off time plus an estimated time enroute, and ending at the actual wheels on time. The demand per quarter hour presented above is also multiplied by 4 to be specified as a rate per hour. In the throughput vs. demand figure the moving average of the throughput is plotted against corresponding demand, with error bars representing one standard deviation in each direction. The window size for the calculation of the moving average is 10 aircraft per hour. The right plot in Figure 5 shows the frequency of the demand in 4 aircraft per hour bins.

It is clear that the throughput increases linearly with demand, until a maximum is approached, at which point throughput does not match demand, but drops off as demand increases. As the number of aircraft waiting increases the throughput of the runway system increases, because more pressure is applied to it. However, there is a point at which the throughput reaches a maximum or saturation value. Beyond this point additional demand pressure (namely congestion or delay in the airspace) is no longer beneficial in terms of increasing throughput. This saturation level is set mainly by the safety separation requirements between aircraft and by controller workload. Therefore, the throughput saturation level is used as an estimate of the maximum service rate capacity of the runway system. As the demand pressure increases further, the throughput is likely to be reduced below the maximum or saturation capacity level, indicating inefficiency. This drop may be due to, among other factors, controller workload constraints, and airspace complexity constraints. When controllers are working too many aircraft they are not able to be as efficient as with fewer aircraft, and the service rate of the resource can thus be reduced. Also complex interactions between flows in the airspace may lead to gridlock as the number of aircraft increases. As the system approaches gridlock the throughput of the affected resource is reduced.

In order to estimate the actual arrival rate capacity of the airport, a hyperbolic curve is fitted to the average throughput with demand less than the drop off point (the dashed line in Figure 5). The second plot in Figure 5 showing the frequency of the demand in 4 aircraft per hour bins shows that the majority of data points fall to the left of this drop-off in throughput. The hyperbolic curve fit asymptotes to throughput equaling demand on the left (45° line) and to a maximum throughput on the right. This maximum throughput represents the actual capacity operated at the airport, for the reported AAR under question. As an example, for a reported AAR of 44 aircraft per hour at Newark, this actual (or asymptotic) arrival capacity is 49 aircraft per hour. Results for other reported AARs and other airports are presented in Table 1 below. Because the hyperbolic curve is fitted to the average arrival throughput, this actual capacity represents an arrival capacity average over a range of other varying factors including, for example, departure rate.

An estimate of the maximum airport arrival capacity achievable can be identified from the raw data as the maximum throughput achieved in any half hour period. Because of binning errors (errors resulting from the use of time bins to estimate a maximum capacity) the FAA recommend that capacity be identified using half hour periods, and not quarter hour periods. Because the maximum is likely to be a rare occurrence, and may represent human error in violating separation requirements, the 99th percentile of the throughput per half hour is thought to be a better estimate of this maximum capacity than the 100th percentile. For a reported AAR of 44 aircraft per hour at Newark, this is 82 aircraft per hour – nearly twice the reported rate. Results for other reported AARs and other airports are presented in Table 1 below.

The capacities corresponding to the AARs reported at PHL, EWR, LGA, JFK, and TEB on the days analyzed in this study are presented in Table 1 below. Three levels are shown: the reported capacity; the asymptotic capacity; and the 99th percentile capacity, which represents the AAR describing the maximum capacity achievable. The charts supporting the data in Table 1, and further details of the capacity analysis performed, are presented in Appendix B.

Table 1. Airport Capacities

Airport	**Reported Capacity (Reported AAR)** [ac/hr]	**Asymptotic Capacity (Actual Configuration Capacity AAR)** [ac/hr]	**99th Percentile Capacity (Max Achievable AAR)** [ac/hr]
PHL	36	41	60
	52	61	72
LGA	31	41	52
	34	41	52
	39	41	52
	42	45	56
EWR	34	40	56
	38	44	56
	40	44	56
	44	49	56
JFK	33	44	44
	35	43	48
	51	60	60
TEB	32	43	28

It should be noted that the 99th percentile capacity is lower than the actual asymptotic capacity for TEB, and very close to it at JFK. This is due to the limited throughput saturation at these two airports under high demand levels, as shown in Figure

6 below for JFK, and explained in detail the next section. The demand at these two airports did not reach levels higher than capacity often enough to cause throughput saturation.

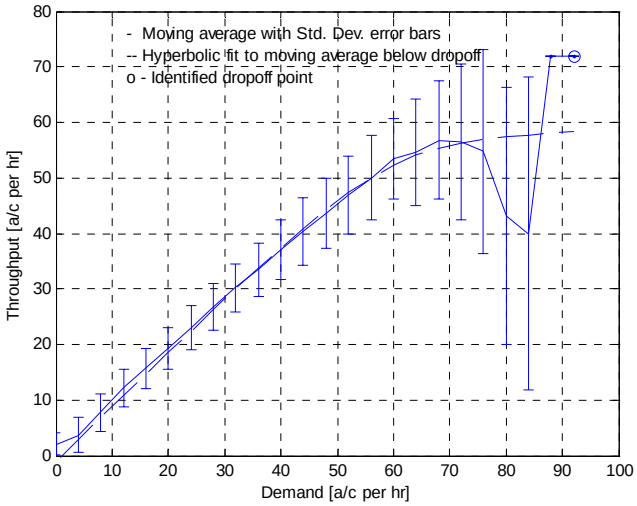

Figure 6. Limited saturation at JFK airport – reported AAR of 51 ac/hr.

The arrival service rate of a runway system is also dependent on the departure service rate when the arrivals and departures share the same or interacting runways. This is demonstrated in Figure 7 where the arrival and departure rates for each period of time are plotted for PHL, when it reported an AAR of 52 aircraft per hour. As the rates increase a tradeoff is evident where serving more arrivals is accomplished at the expense of serving less departures and vice versa, resulting in a capacity envelope.

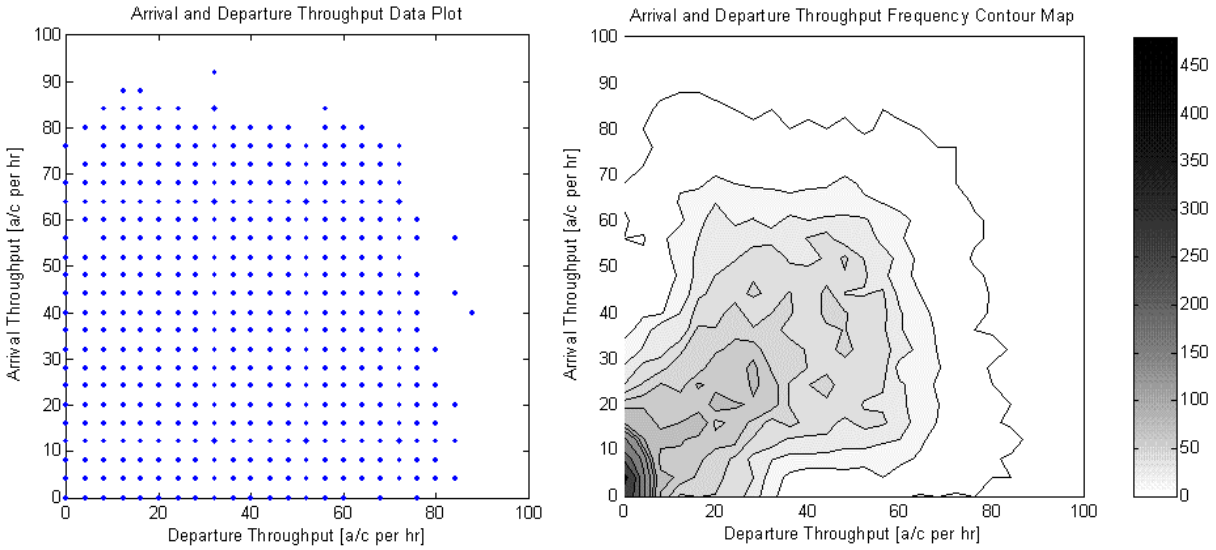

Figure 7. Capacity envelope for PHL with a reported AAR of 52 ac/hr; June, July and August 2001.

This capacity envelope becomes an important parameter, as any increase in the system throughput should be realized through an increase in the capacity or utilization of the runways. An outward shift of the envelope represents an increase in available capacity. In some cases an airport may be operating consistently below the envelope and an increase in utilization of available capacity represents moving the operating point closer to the envelope. It is also important for modeling the tradeoff between arrival and departure maximum rates. For example, if the departure rate for a given period is known and fixed, the maximum arrival rate may be read off the envelope and applied to the simulation.

The effect of departures was thus modeled by calculating a number of airport capacities in the same way as in Figure 5, for each reported AAR, for a range of actual departure throughputs. Actual departure throughput was then plotted against the calculated arrival capacity for that departure throughput, yielding a capacity envelope for that reported AAR. Capacity envelopes were thus developed for the asymptotic capacity and a range of percentile capacities from the 85^{th} percentile to the 99^{th} percentile at EWR and LGA, and from the 90^{th} percentile to the 99^{th} percentile at PHL, JFK and TEB. All percentiles of the capacity envelopes were calculated per quarter hour. The 99^{th} percentile capacity envelope was also calculated per half hour, and it is this envelope that represents the maximum achievable capacity at the airport.

Because of the discrete increments between percentiles in the capacity envelopes calculated, a number of the envelopes for different percentiles were found to overlap, and thus not model any change in capacity with the change in percentile. This produced inaccurate estimates of capacity at the different percentiles. The percentiles were thus estimated by fitting a gamma distribution to the throughput data, and calculating the percentile from the fitted gamma distribution instead of the original throughput data. Confidence intervals of 90% were also calculated for the parameters of the gamma distribution, and if the parameters calculated did not fall within these intervals the fitted distribution was assumed to be too poor, and discarded. In these cases the percentiles were calculated from the original throughput data instead.

As described above capacity envelopes were generated for each reported AAR at each airport. This assumes that for each reported AAR there is only one reported airport departure rate (ADR). If ADR varied with constant AAR a different capacity envelope could be developed for each AAR/ADR combination. Such cases were very limited over the 15 days from November 2003 analyzed in this study. In almost all cases there was only one ADR for any given AAR. However, there were a few cases at PHL and JFK where ADR varied with AAR. These cases were thus dealt with independently, and capacity envelopes were developed for each AAR/ADR combination in each case.

The capacity envelopes are shown in Figure 8 to Figure 12 below for a single commonly reported AAR at each airport. Estimating departure rate for a given period as the actual departure throughput for that period (from the ASPM database per quarter hour), the asymptotic arrival rates and arrival rates at each percentile were thus read off the appropriate envelopes and applied to the simulation.

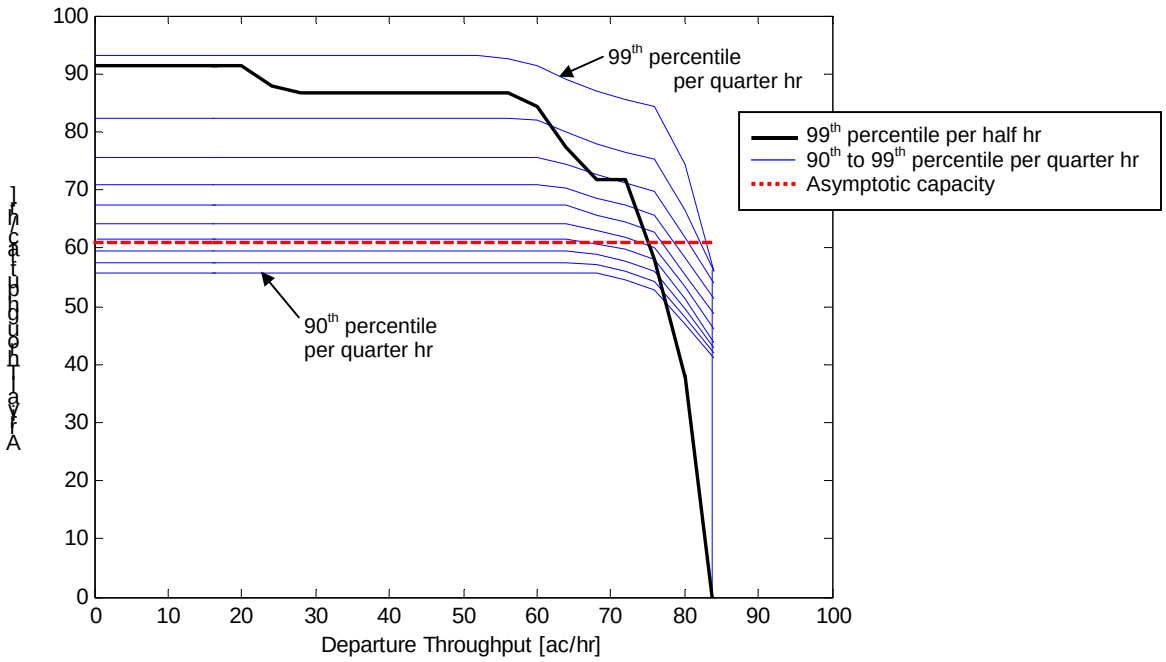

Figure 8. Capacity envelope for PHL with a reported AAR of 52 ac/hr, showing the asymptotic capacity envelope (thick dashed red), modeled capacity envelopes for the 90th to 99th percentiles capacities per quarter hour (thin blue), and the 99th percentile capacity envelope per half hour (thick black).

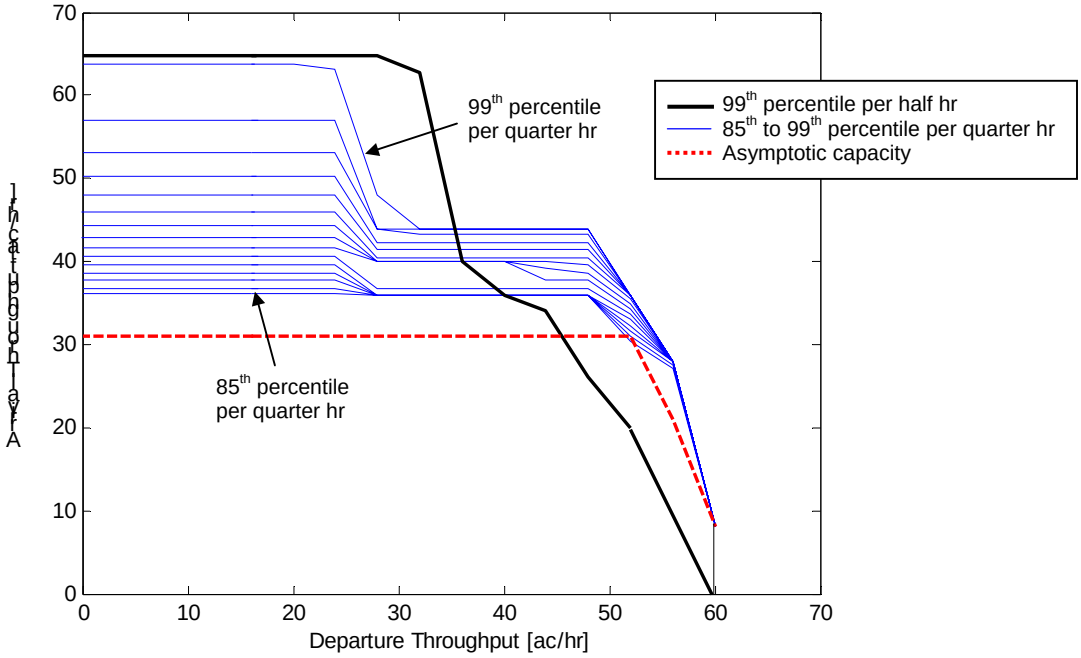

Figure 9. Capacity envelope for LGA with a reported AAR of 24 ac/hr, showing the asymptotic capacity envelope (thick dashed red), modeled capacity envelopes for the 85th to 99th percentiles capacities per quarter hour (thin blue), and the 99th percentile capacity envelope per half hour (thick black).

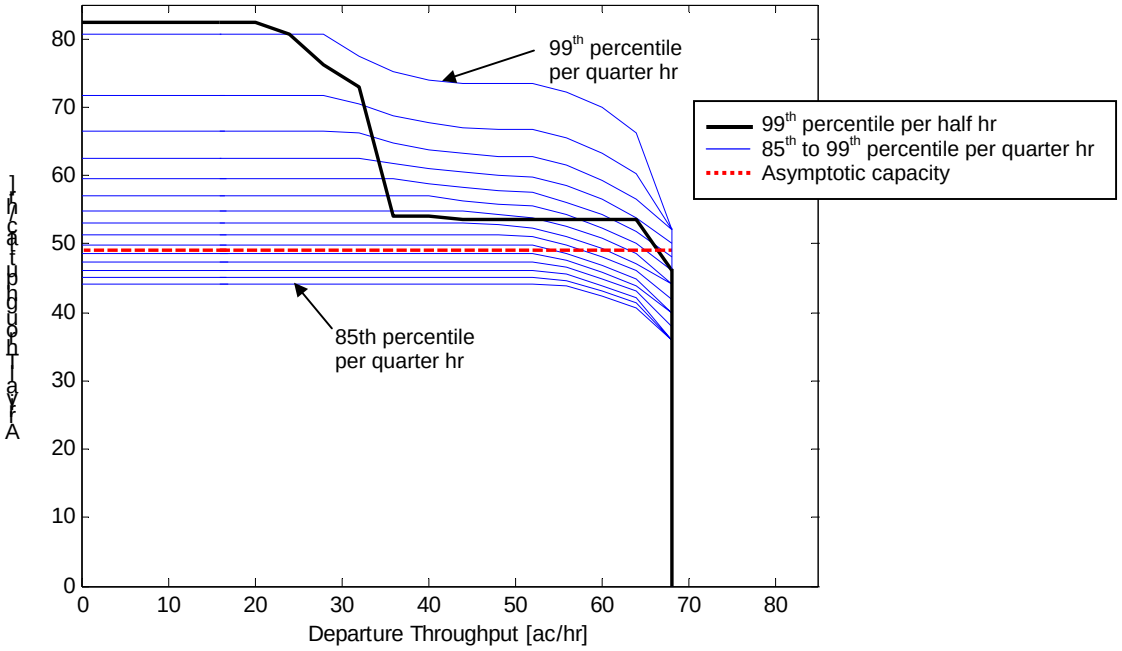

Figure 10. Capacity envelope for EWR with a reported AAR of 44 ac/hr, showing the asymptotic capacity envelope (thick dashed red), modeled capacity envelopes for the 85th to 99th percentiles capacities per quarter hour (thin blue), and the 99th percentile capacity envelope per half hour (thick black).

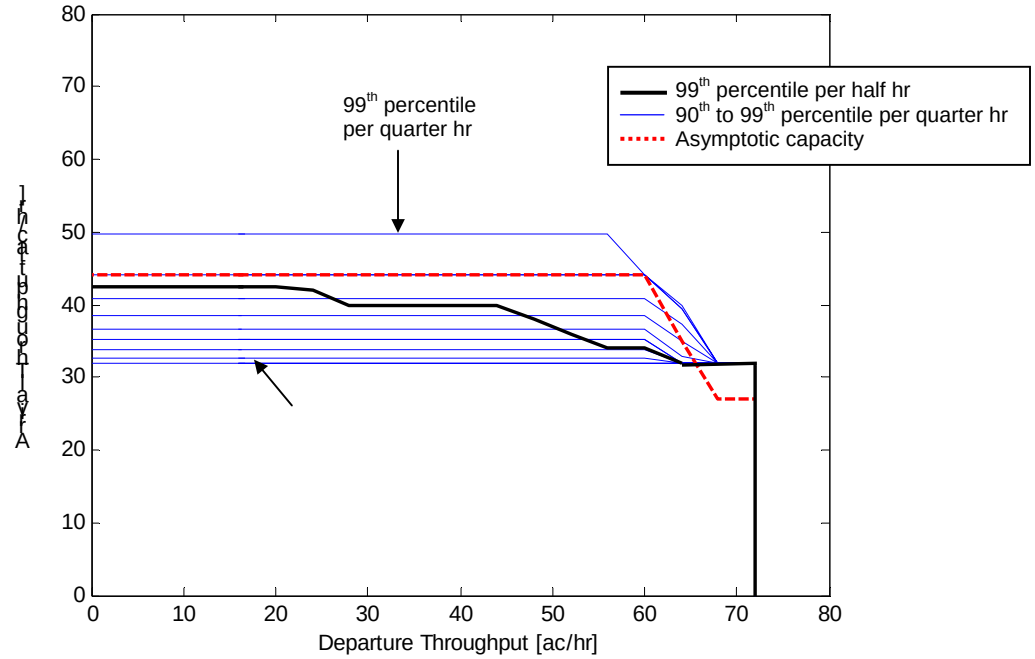

Figure 11. Capacity envelope for JFK with a reported AAR of 32 ac/hr, showing the asymptotic capacity envelope (thick dashed red), modeled capacity envelopes for the 90th to 99th percentiles capacities per quarter hour (thin blue), and the 99th percentile capacity envelope per half hour (thick black).

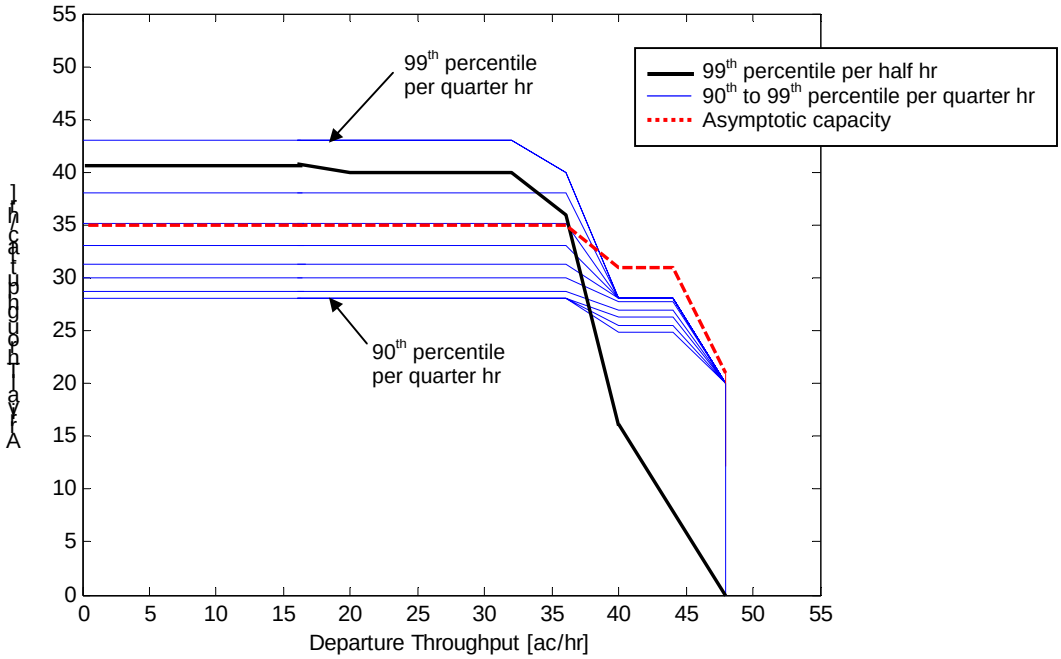

Figure 12. Capacity envelope for TEB with a reported AAR of 28 ac/hr, showing the asymptotic capacity envelope (thick dashed red), modeled capacity envelopes for the 90th to 99th percentiles capacities per quarter hour (blue), and the 99th percentile capacity envelope per half hour (black).

It is clear from the capacity envelopes presented that the 99th percentile envelope calculated per half hour is significantly different to that calculated per quarter hour, particularly at the higher departure rates. This suggests that the binning error is significant, at least under certain circumstances. Many of the other percentile envelopes, calculated per quarter hour, are higher than the 99th percentile envelope calculated per half hour at the higher departure rates. This does not mean that these envelopes cannot be applied, however, because the 99th percentile envelope calculated per half hour is only an estimate of airport capacity. As long as the modeled airport throughput under the applied percentile calculated per quarter hour does not violate this limit more than occurs under actual operations, the maximum airport capacity is not violated. The method by which the airport capacity to be applied to McTMA is identified is described in detail in Section 7.2.

The asymptotic capacity envelope is well below than the 99th percentile envelope in all cases except JFK, where it is higher, and TEB, where it is only just lower. This is consistent with the analysis presented in Table 1, suggesting that these airports do not saturate to the extent of the other airports.

3.1.2. TRACON Capacity Constraints

According to previous NASA study on Philadelphia [7] the bottleneck in the flow to PHL are the runways at Philadelphia airport. The primary constraint is therefore the airport acceptance rate. According to interviews with TMCs at N90, the airspace within N90 is the primary constraint to arrival flows into N90, and not the arrival rates at the airports. This is due primarily to the large number of airports (3 major and 12 satellite airports) that share a relatively small airspace. The interaction of flows, the number of aircraft within the TRACON, and the size of the TRACON, all limit flow into each airport.

JFK for example does not generally run at capacity, particularly when the reported AAR is high (such as 51 aircraft per hour). According to interviews with N90 TMCs during the site visit, JFK operates three arrival runways during the international arrival push from 1pm to 5pm. This is because international flights take priority over domestic flights, as they are generally fuel critical. However, during other hours JFK is not able to operate three arrival runways, and consequently operates at a significantly lower throughput. This is because of the flows into the other three primary airports in N90. These observations from the N90 TMCS are confirmed in Figure 6, which shows limited saturation of JFK throughput as demand increases and when the reported AAR is high, unlike the other airports. TEB displays an almost complete lack of throughput saturation.

Another indication of the N90 airspace constraints is the ongoing efforts to improve the airspace patterns and reduce the interaction between the flows. An example of such an improvement is the flipping of the arrival fixes into LaGuardia and Newark airports, from ZDC, which was implemented in the summer of 2002. The locations of these fixes used to require LaGuardia and Newark arrival streams to cross. After the fixes were flipped the flows no longer cross. According to TMCs at ZDC and N90 this has improved operations significantly and allowed holding 2 aircraft in the TRACON, while no holding was exercised before.

In order to compare the individual airport versus TRACON capacity limitations, Figure 13 and Figure 14 plot the actual arrival rate versus the reported arrival rate for each of the three major N90 airports separately and for all airports combined (representing the TRACON as a single resource). While each airport exhibited a certain degree of underutilization of the runway arrival capacity (measured relative to the reported capacity), the underutilization was more pronounced for the combined airports. This indicated that when some of the airports operate at capacity the other airports usually operate below their capacity. This observation supports the N90 personnel comments made during the site visit that the N90 airspace is a more binding constraint than the N90 airports runway capacity.

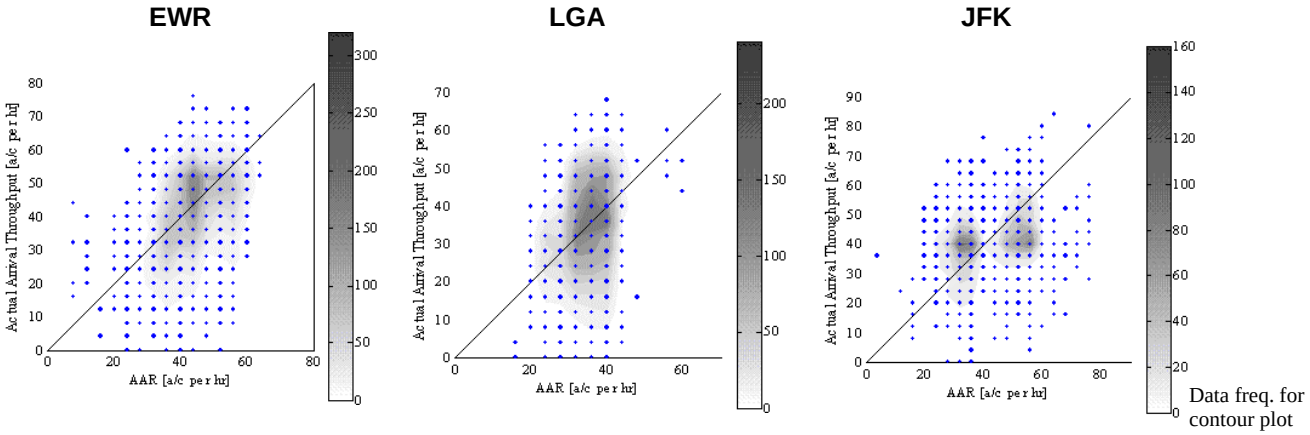

Figure 13. Utilization of N90 airports.

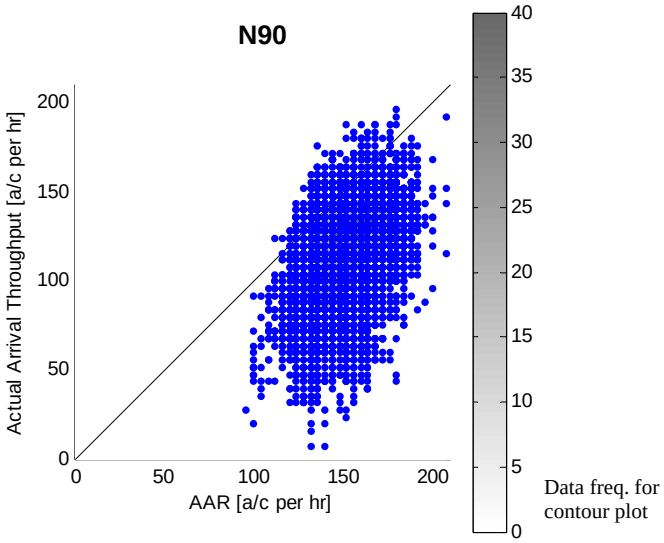

Figure 14. Utilization of N90 TRACON.

One implication of the observation that the N90 airspace is the bottleneck for the N90 arrival flows is that an increase in the N90 acceptance rate (due to McTMA for example) may easily be taken advantage of by utilizing existing runway capacity. It was indicated by a TMC during the site visit, for example, that JFK airport has runway capacity that is unused due to TRACON airspace limitations (sharing the TRACON with other airports). This is unlike other locations where the limiting factor is the available runways and adding more runways is needed in order to increase capacity.

Figure 13 shows a certain level of inefficiency and underutilization of capacity in the current operations. The benefits of McTMA in increasing the system throughput may be realized through an increase in the utilization of the available capacity or in the available capacity as will be described in Section 5. One implication of the TRACON being the capacity limiting resource is to use the combined airport throughput versus

demand curve to estimate the system capacity instead of using the individual airport capacity models. This analysis was completed for a number of runway configuration combinations at the primary airports within N90 (EWR, LGA, JFK, and TEB). The throughput versus demand curves are of the same shape as those generated for the individual airports, as shown in Figure 15. A hyperbolic curve was thus fitted to the data to the left of the drop-off in throughput, as for the individual airports, and the horizontal asymptote identified as the actual arrival capacity of the TRACON.

The TRACON capacity limit was not imposed in the modeling either McTMA or baseline operations. This was due to the immaturity of the N90 capacity analysis, particularly with inclusion of the departure rates and modeling of the capacity envelopes for each airport. McTMA also applies airport acceptance rates and not a TRACON acceptance rate. It is unclear how a TRACON constraint will be applied in McTMA operations.

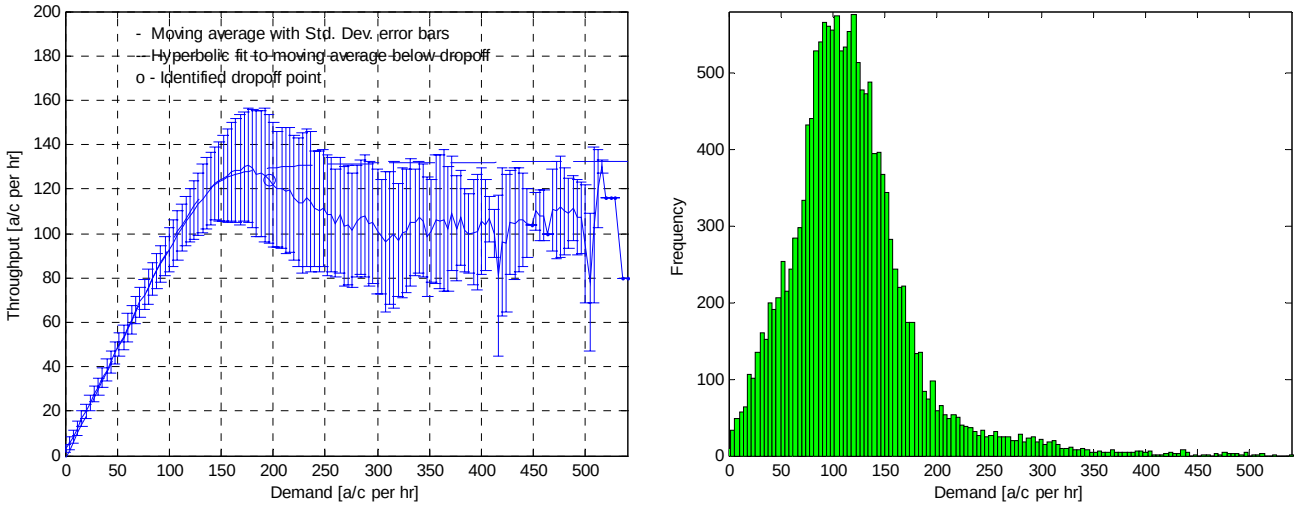

a) b)

Figure 15. TRACON throughput plotted against demand – a) moving average plot with hyperbolic curve fit, and b) demand frequency plot. Data source: ASPM database for January to August of 2001, and June to August of 2000.

Under current operations the TRACON airspace limitation often induces the N90 TMCs to impose MIT restrictions due to airspace congestion as opposed to runway capacity limitation. The main tool that is used to set the acceptance rate of N90 is Miles In Trail, which are passed back to the adjoining ARTCCs.

MIT restrictions at the PHL TRACON boundary are primarily specified according to airport acceptance rate. However, according to expert elicitation from N90 TMCs, MIT restrictions at the N90 TRACON boundary are specified according to required airport acceptance rates, limits on the number of aircraft in the N90 airspace, demand on the TRACON, and local delays.

The modeling of current operations, and the specification of MIT according to airport capacities is presented in detail in Section 4.

3.1.3. Sector Capacity Constraints

While the capacity of sectors to absorb delays and hold aircraft is usually not the constraining element in the flow, it is an important parameter for the propagation of delays upstream. When restrictions are imposed because of a runway or TRACON limited acceptance rate, the delays are propagated back to the upstream sectors. Each sector has a limited ability to absorb delay and hold aircraft and once this limit is reached the delays are propagated further upstream. It is therefore, essential to identify the sector capacities in order to model the current baseline operations and the McTMA metering process, both of which depend on this parameter.

According to expert elicitation from TMCs during the site visits, a number of factors affect the capacity of sectors under normal weather conditions:

1. **Compression due to descent** – As arrivals descend, they are generally required to slow down. If a sector descends aircraft, more spacing is required at the entry point in the sector to provide the 5 MIT spacing (or more if the flow is restricted) required at the exit point in the sector, because of the speed reduction of the leading aircraft as it descends. These sectors thus have reduced capacity as they are not able to absorb as much delay. This particularly affects sectors close to the TRACON.
2. **Vectoring** – The proximity of other airways can limit the amount of vectoring aircraft are able to perform because the aircraft are limited in how much they deviate from their airways without violating safety separation requirements with other traffic. This is particularly a problem in ZNY as there are number of close parallel airways between N90 and ZOB. As a result ZNY exercises very limited vectoring. There is more space for such vectoring in ZBW, ZOB and ZDC.
3. **Speed Reduction** – Only a small amount of delay is able to be absorbed through speed reduction depending on the size of the sector, as speed can generally only be reduced by a few knots. The amount of speed reduction possible can be greatly affected by wind. If the prevailing wind is a tail wind, the effect of speed reduction is limited as aircraft fly faster with respect to ground, while any kind of headwind can be used very effectively to open gaps in the flow using speed reduction. The prevailing wind is generally from the west, so speed reduction for N90 and PHL flows in ZNY and ZOB has limited effect. Speed reduction for N90 and PHL flows in ZBW, however, can be very effective. Speed reduction for N90 and PHL flows in ZDC is affected by wind to a lesser degree as the prevailing wind is a crosswind, and often not a tailwind or headwind.
4. **Holding** – In extreme cases when aircraft must be delayed by an amount that is too great to be absorbed by vectoring or speed control, the aircraft must be put into a holding pattern. Holding patterns are limited in location and capacity by other airways, alongside and overhead. The minimum and maximum altitudes of an airway define the number of aircraft that can be held in a holding pattern on the airway, as the vertical separation between aircraft must be 1000ft. Because the sectors in ZNY are highly constrained by the number and proximity of the parallel airways between N90 and ZOB, ZNY does not hold aircraft. ZOB however, is able to hold aircraft bound for N90 on the boundary of ZNY. There is also substantial holding capacity in both ZBW and ZDC.

5. **Gridlock** – The possibility of gridlock limits the ability of a sector to absorb delay. One interesting observation at ZOB is that the holding patterns of the N90 arrivals heading to ZNY can block the N90 departure routes from ZNY into ZOB. This limits ZOB's ability to hold ZNY arrivals and forces ZOB to holds the N90 departures from ZNY when ZNY holds the N90 arrivals from ZOB. Because the arrival and departures are interdependent, where departures need to leave the airports in order to make room for the arrivals to land, a gridlock effect is created where both arrivals and departures are holding and waiting each for the other to advance.
6. **Complexity of flow patterns** – The degree of interaction between the flows in a sector can affect the capacity of the sector. If flows descend, climb, or cross, the capacity can be greatly reduced due to the high workload associated with controlling the complex patterns. According to expert elicitation from ZOB TMCs, this is particularly true for example in Lorain sector in ZOB, in which no delay can be absorbed. OALTs (Operational Acceptable Level of Traffic) are correspondingly particularly low for such sectors.

The capacity can thus be specified for each sector by a number of variables. These are the amount of delay that can be absorbed in the sector without holding (or sector delayability), the number of aircraft that can be held in the sector, and the OALT for the sector. Sector delayability was determined for some sectors through expert elicitation during the site visits, particularly in ZOB as shown in Figure 16. As shown the maximum delay that can be absorbed without holding is 3 minutes, and as was also indicated during the site visits, delays of about 4 minutes and above are usually absorbed through holding. However, McTMA researchers also indicated that other facilities such as DFW use 6 minutes as a rule of thumb for the duration of a single holding spin. The difference may be caused by the restricted size of holding patterns in the Northeast. Since no explicit analysis was conducted to determine the minimum holding delay a value of 5 minute delay was used in this study to indicate that an aircraft was most likely held. ZNY indicated that they have very little ability to absorb delay through vectoring or speed reduction (due to tail winds), and they have very little holding capacity.

The delay parameters that were identified through interviews on the site visits were confirmed through analysis of Host traffic data as described below.

The delayability between each meter fix pair in the system was also estimated by calculating actual historic transitions times between meter fixes, and comparing these to unimpeded times between the same fixes. The derivation of unimpeded transition times is described in Demand Modeling in Section 3.2. The difference between the actual transition times and the unimpeded transition time represents the amount of delay that was absorbed. Held aircraft were excluded from the analysis by excluding all flights with transition times more than five minutes longer than the unimpeded transition time. Five minutes was chosen, as an aircraft generally requires four minutes to make one spin in a holding pattern. The maximum amount of delay able to be absorbed between the meter fixes, or the delayability, was identified as the 90^{th} percentile of the historical flight delays between the fixes. The upper 10% of delays were assumed to be abnormalities. The sensitivity of the final benefits to delayability is analyzed in Section 9.2.

The delayabilities calculated for the meter fix pairs in ZOB are presented in Figure 16, alongside those identified from expert elicitation. Because some of the meter fix pairs include more than one sector, some of the delayabilities identified from expert elicitation for each sector must be added for comparison to the delayability calculated from historical transition times. It is clear from this comparison that the delayabilities calculated from historical transition times match closely to those identified through expert elicitation, for the downstream sectors. However, for the upstream sectors the delayabilities identified through expert elicitation are significantly higher than those calculated from historical transition times. This is because the upstream sectors are not often required to absorb as much delay as they are able to. The analysis of the historical transition times reflects this. The downstream meter fix pair delayabilities are thus estimated according to the analysis of the historical transition times, but the upstream meter fix pair delayabilities are increased according to the delayabilities identified through expert elicitation. Because delayabilities could only be identified through expert elicitation for ZOB, these numbers were extrapolated to all upstream sectors in the system.

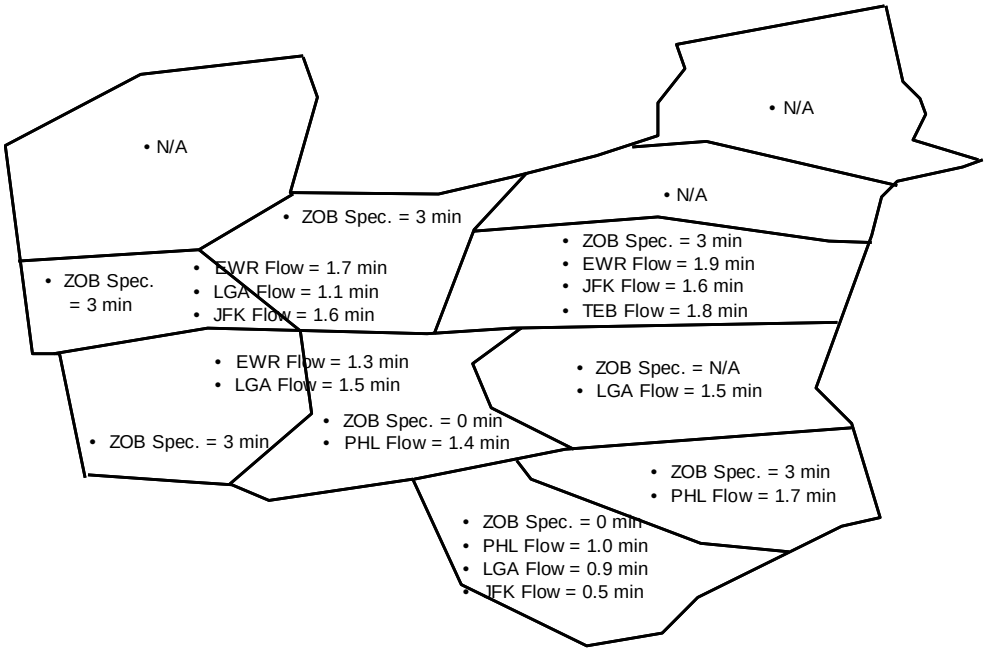

Figure 16. Delayabilities, ZOB.

While the delay that is able to be absorbed in the airspace between meter fixes may be large, the throughput at the downstream resources (runways and TRACON) may saturate at a lower level of delay. This lower level of delay corresponds to the level of demand at which the throughput saturates in Figure 5. It would therefore not be beneficial in terms of throughput to delay aircraft beyond this level (even if the sectors capacity allows for more delay). This delay level at which throughput saturates has not been modeled in this study, as the delayabilities specified in McTMA are not currently to include this effect. Inclusion of the effect may however increase the benefits of McTMA, if applied in the future, and may thus be studied further in the future work.

3.2. Traffic Demand

The demand on the system is driven by the airlines scheduling of flights into the airports served by the TRACON. The demand at a fix is represented by a series of Estimated Times of Arrival (ETA) calculated based on the airline schedules and flight plans. If unimpeded, aircraft fly from fix to fix at speed. Using unimpeded transition times between fixes, estimated times of arrival (ETAs) can be calculated at each fix, and at the airport as shown in Figure 17.

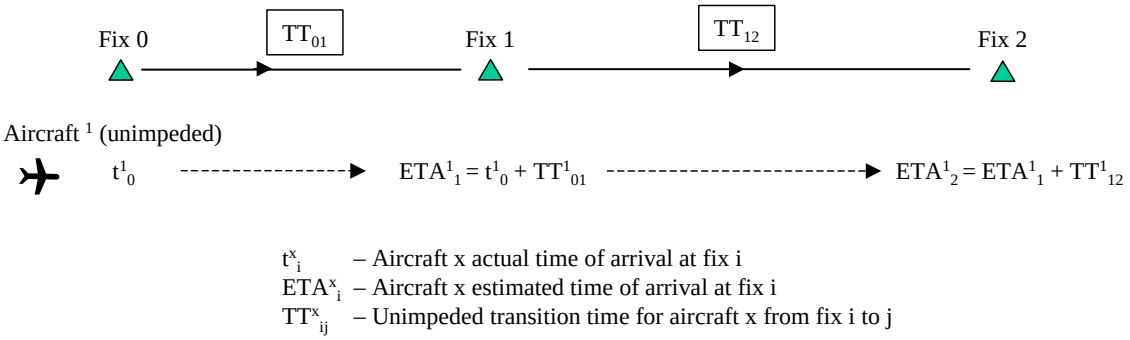

Figure 17. Demand based on estimated times of arrival

ETAs at fixes are calculated according to a flight's initial conditions, including time of entry into the system (t_0). This entry time into the system corresponds to the actual time that the aircraft crosses the outmost perimeter of the system. Other inputs include unimpeded transition times to the downstream fix in the flight plan (TTs), and unimpeded transition times between subsequent fixes. Equations (3) describing the subsequent calculation of ETAs are as follows:

$$\begin{aligned} ETA_1 &= t_0 + TT_{01} \\ ETA_2 &= ETA_1 + TT_{12} \ldots \\ &or \\ ETA_j &= ETA_i + TT_{ij} \end{aligned} \qquad (3)$$

In this manner each flight's ETAs are calculated for all the applicable points, working downstream from the system boundary to the runway threshold.

Figure 18 describes an arrival flow network for JFK. The network includes 3 tiers, each presented in different colors. The flows were identified according to STARS, flows illustrated in presentations by Cleveland ARTCC Traffic Management Unit and Boston ARTCC, and according to host track data from September 12, September 17, and September 19, 2002. The arrows in the figure represent the flows modeled. The flows are presented for all other airports under consideration in Appendix A, along with further details on how they were generated.

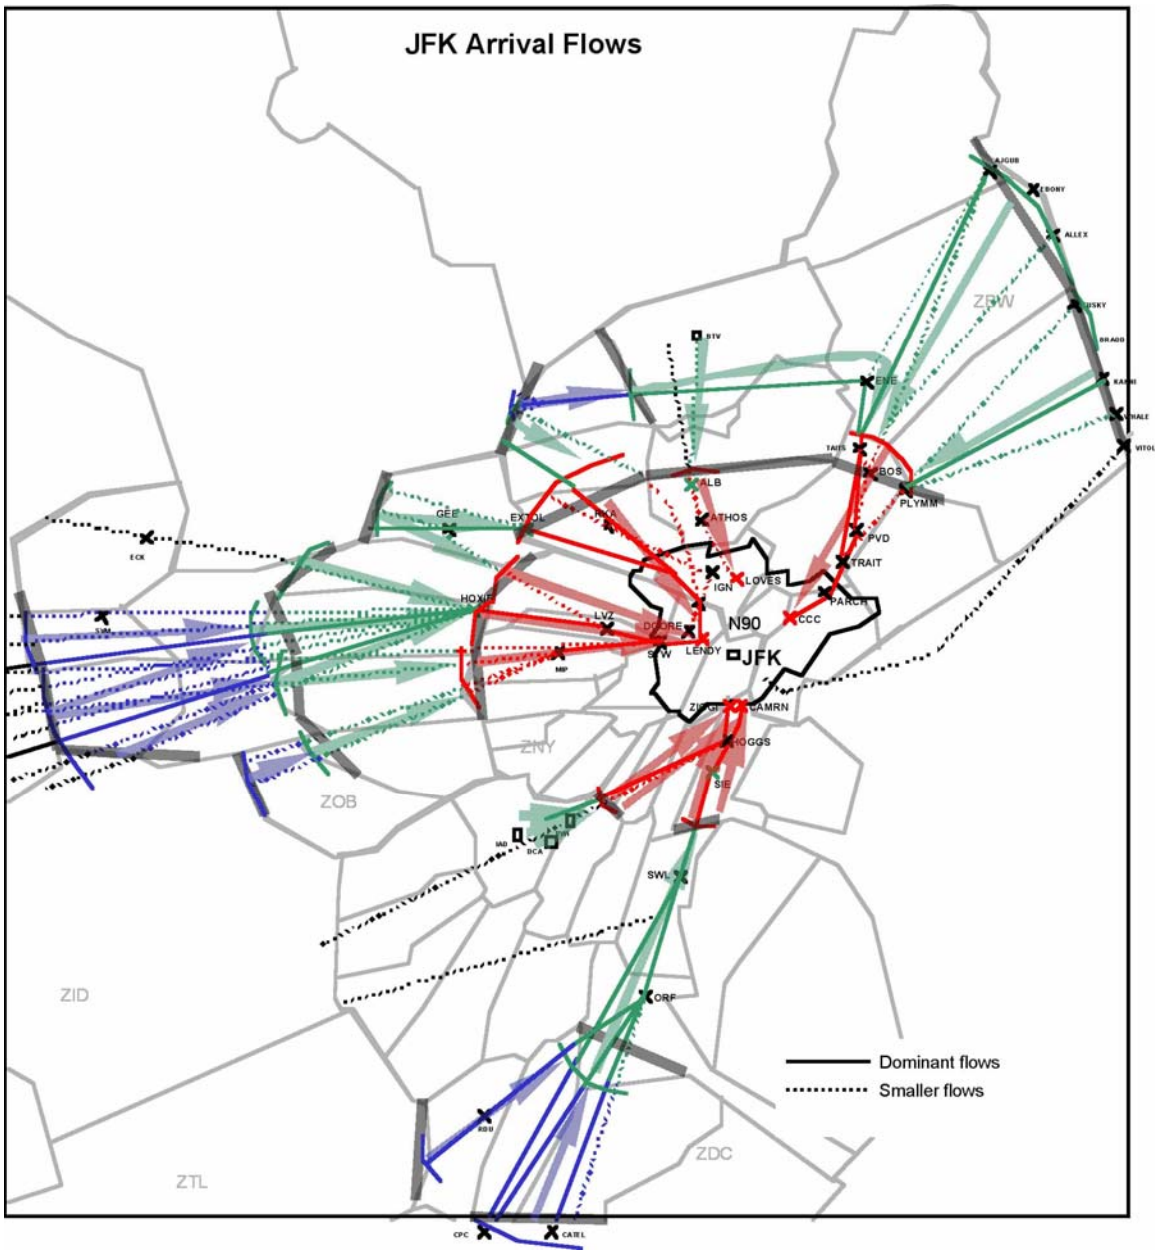

Figure 18. Arrival flow network and multiple tiers for JFK

Meter fixes on the flows in Figure 18 were chosen according to an approach suggested by NASA McTMA researchers. This approach includes specification of meter fixes and meter fix freeze horizons as close to sector boundaries as possible, but also at flow merge points. Meter fix arcs were also specified in some cases, instead of meter fix points, to ensure that as much traffic as possible is metered.

The flight locations at the start of the simulation were extracted from actual host track data. In the case of internal departures for airports that had a significant amount of traffic the entry time into the system was taken as the recorded first track point after departure. Internal departures from smaller airports, which did not contribute

significantly to the traffic, were captured as they crossed the next tier or the TRACON. An airport was considered to have a significant amount of traffic if traffic from this airport constituted more than approximately 1% of arrival traffic into the destination airport under consideration. This generally includes any airport with more than 5 flights a day to the destination airport under consideration.

Demand is based on actual entry times into the system, as opposed to scheduled entry time, for both airborne and internal departure aircraft. The resulting demand is conservative because McTMA is expected to increase demand into the system by reducing the need and severity of restrictions imposed upstream of the system. For example, because of McTMA, internal departures may be delayed less on the ground than in current operations, and airborne aircraft may be delayed less upstream, due to GDPs and upstream MIT, than in current operations.

Unimpeded transition times were estimated according to a statistical analysis of actual historic transition times for flights passing through each meter fix pair. The actual transition times for all flights passing through each meter fix pair from fifteen days in November 2003 (Nov. 8, 9, 11, 12, 14, 19, 20, 22, 23, 24, 25, 26, 27, 28, and 29) were plotted against the queue size encountered, for 2 different weight classes, for 5 different wind conditions, and by runway configuration (transition times within the TRACON only). Actual transition times and queue sizes were calculated from host track data for these days. The queue size experienced by a flight between each pair of fixes was calculated as the number of aircraft that passed through the downstream fix from the time when the flight under question crossed the upstream fix, to when it crossed the downstream fix. The weight classes plotted were Small and Other (including Large, B757, and Heavy). Transition times were not separated for Large, B757 and Heavy weight classes because the resulting unimpeded transition times were not found to vary significantly over these weight classes. The wind conditions were separated by plotting hourly RUC wind speed and wind angle at 30,000ft at the center of ZNY for the 15 days studied, and identifying dominant clusters, as shown in Figure 19. The wind clusters identified were for a northerly wind (wind angle greater than 90º), a strong westerly wind (wind angle between 45º and 90º, and a wind speed greater than 150kts), a weak westerly wind (wind angle between 45º and 90º, and a wind speed less than 150kts), a strong southerly wind (wind angle less than 45º, and a wind speed greater than 150kts), and a weak southerly wind (wind angle less than 45º, and a wind speed less than 150kts). The northerly wind cluster was not separated into weak and strong clusters because of the lower number of data points in this cluster. The resulting unimpeded transition times were found to vary significantly (in the order of a few minutes in some cases) by wind conditions. The RUC data was obtained from NASA.

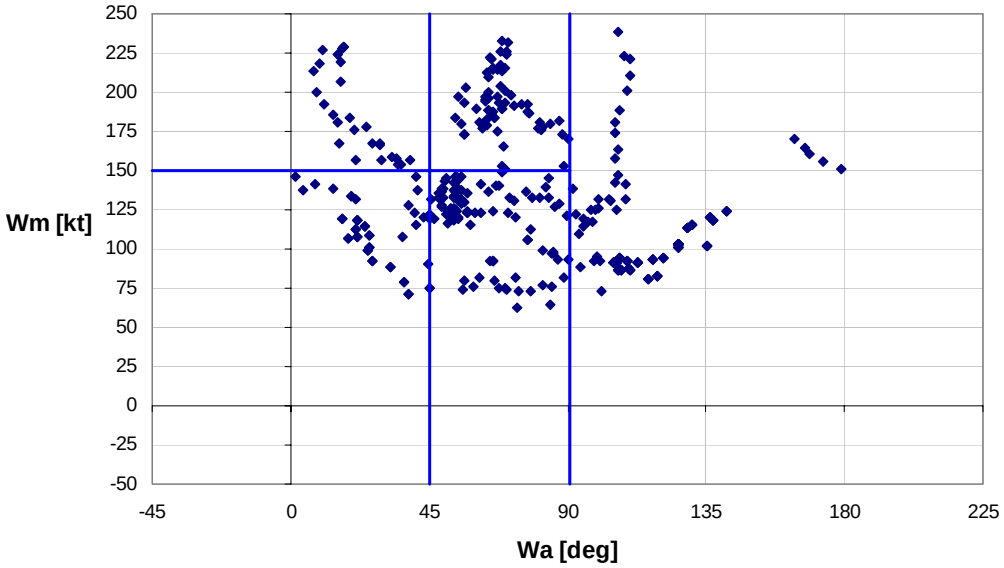

Figure 19. RUC wind speed (Wm) and wind angle (Wa), with dominant clusters separated.

Figure 20 below shows the queuing dynamics for the flow from the boundary of ZNY and ZOB to PENNS – the west arrival fix into EWR, for Large aircraft, and a weak westerly wind. Similarly Figure 21 shows the queuing dynamics for the flow from PENNS to EWR, landing on runway 22L or 22R, for Large aircraft, and a weak westerly wind. Such data was generated for each meter fix pair in the system.

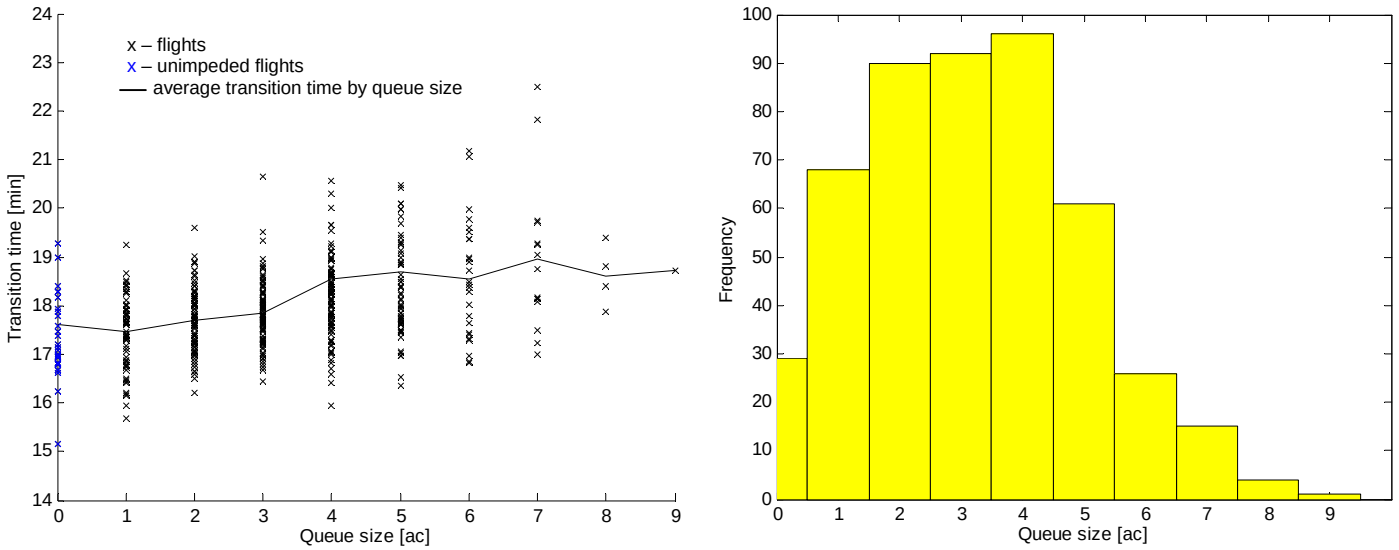

a) b)

Figure 20. a) Queuing model for Large weight class, under a weak westerly wind (wind angle between 45º and 90º, and a wind speed less than 150kts) from the boundary between ZNY and ZOB to PENNS. b) Corresponding frequency distribution with queue size.

x – flights
x – unimpeded flig...
average transition time by queue size
parabolic curve fit

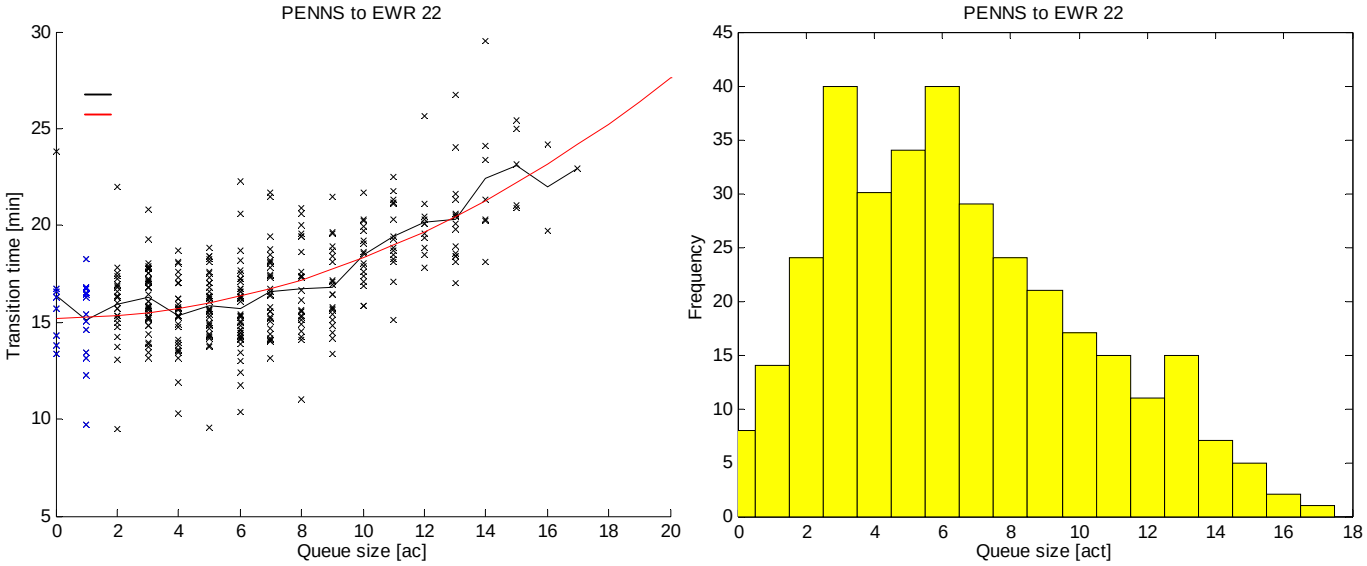

a) b)

Figure 21. a) Queuing model for Large weight class, under a weak westerly wind (wind angle between 45º and 90º, and a wind speed less than 150kts) from PENNS to Newark airport, landing on runway 22L or 22R. b) Corresponding frequency distribution with queue size.

It can be seen that transition times increase with increasing queue size in both Figure 20 and Figure 21. The increase in transition time in Figure 20 is slow, and approximately linear, while that in Figure 21 is faster and of a higher order than one. The figures thus suggest that there is more of a queuing effect in the TRACON (Figure 21) than outside the TRACON (Figure 20). Similar plots for meter fix pairs further upstream also show less of a queuing effect than in Figure 21. The TRACON is expected to exhibit more queuing than upstream because it is more constrained and more aircraft are required to queue in a smaller region of airspace. This is because many streams from upstream join in the TRACON to form the final approach queue to the runway. Upstream, there are fewer aircraft and more airspace.

Unimpeded transition times upstream of the TRACON were estimated by fitting a normal distribution to the data points with low queue size, and sampling from this distribution. This introduces variability to the estimated unimpeded transition times, which models the different unimpeded transition times that result for different flights plans and aircraft types flying through the meter fixes modeled. The low queue size threshold was identified as the lowest queue size for which there were at least 10 data point. This ensures that enough data points are identified for fitting a distribution. This includes all data with blue x's in Figure 20 above. In this case there were enough data points with a queue size of zero from which to generate the distribution. Held aircraft were excluded from the analysis by excluding all flights with transition times more than five minutes longer than the average transition time. Five minutes was chosen as to represent held aircraft as in the calculation of delayabilities in Section 3.1.3.

Because of the high traffic within the TRACON, other than exhibiting more queuing, there are also generally fewer cases of low queue size than upstream. This means that an estimate of the unimpeded transition time using the 10 data points with

lowest queue size may include data points with queue sizes as high as 4 aircraft, which may have transition times significantly higher than those with a queue size of zero aircraft. Such an estimate of unimpeded transition time may thus be unrealistically high. For this reason unimpeded transition times within the TRACON were estimated by fitting a curve to the data, and estimating the average unimpeded transition time as the zero queue intercept of this curve fit. A normal distribution was fitted around this average with a standard deviation equal to that of the data points with low queue size. Two curves were fitted to the data in each case – a 2nd order parabolic curve fit, and an exponential curve fit. The quality of the fit was compared by calculating the sum of the square of the residual for each curve fit. The curve fit with the lower value for this parameter was chosen to model the queuing effect. In the case presented in Figure 21 the 2nd order parabolic curve fit fitted the data more accurately than the exponential curve fit, and was thus chosen to model the queuing effect. In other cases, such as that presented for the flow from BUNTS to PHL airport, landing on 27L or 27R, shown in Figure 22, the exponential fit was found to model the effect more accurately.

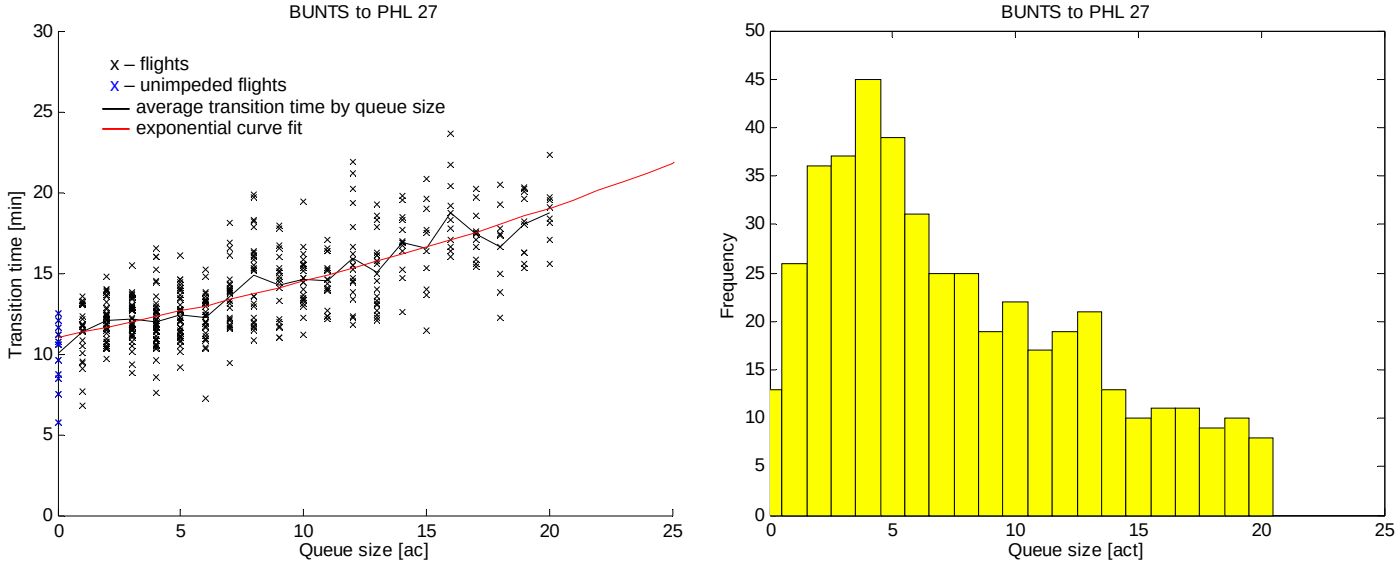

a) b)

Figure 22. a) Queuing model for Large weight class, under a weak westerly wind (wind angle between 45º and 90º, and a wind speed less than 150kts) from BUNTS to Philadelphia airport, landing on runway 27L or 27R. b) Corresponding frequency distribution with queue size.

ETAs were calculated according to the unimpeded transition times sampled from the distributions generated, and according to aircraft weight class and wind conditions, for all the applicable points for each flight, working downstream from the system boundary (the freeze horizon of the third tier) to the runway threshold. This estimated time of arrival based on unimpeded travel between fixes represents the baseline relative to which en-route delay is accumulated.

3.2.1. Comparison of Statistical Model to Trajectory based Model

The results of the statistical model developed to estimate unimpeded transition times was compared to the results of a trajectory based model, the CTAS Trajectory Synthesizer (TS). Given initial conditions, flight plans, RUC wind files, and aircraft and engine type data the TS calculates a flight's ETA by modeling its trajectory explicitly. A number of flights from November 20, 22, and 26, 2003, were analyzed using the TS, and their ETAs at the TRACON boundary compared to those calculated using the statistical models described above. Flight ETAs were not calculated at the runway because the adaptation for the TS did not include TRACON internal routes. The differences between the ETAs are presented in the histograms in Figure 23, for each airport.

In each of the histograms presented in Figure 23 there is a peak at zero. This suggests that for a number of flights there is little difference between the ETAs calculated at the TRACON using the TS, and using the statistical models presented above. In each case, however, there is a tail on the positive side. This tail is particularly large at LGA and JFK. Positive differences indicate that the TS modeled ETA for these flights is later than that modeled using the statistical model presented above. The average and median of the difference in ETA is presented in Table 2 below.

Table 2. Mean and Median of Differences between TS and Statistically modeled TRACON ETA.

Airport	Number if Data Points [Flights]	Mean Diff. between TS and Stat. TRACON ETA [min]	Median of Diff. between TS and Stat. TRACON ETA [min]
PHL	782	0.806	0.310
LGA	595	2.385	0.264
EWR	562	1.698	0.259
JFK	282	2.472	1.593
TEB	279	1.423	0.000

It is clear from Table 2 that the tail on the positive side of the histograms in Figure 23 has an effect on the mean differences between the TS and statistically modeled ETAs, as each mean value is greater than zero. For LGA and JFK the mean difference is particularly high, and is in the order of magnitude of delay incurred in the system. The medians, however, are not as high, except at JFK, and show good correlation between the TS and statistical model. A median is affected less by a tail, and is thus expected to show better correlation in this case because the tail appears to be the primary cause of the high differences. It is important, however, to identify the causes of the tail, and the cause of the high median of the differences at JFK.

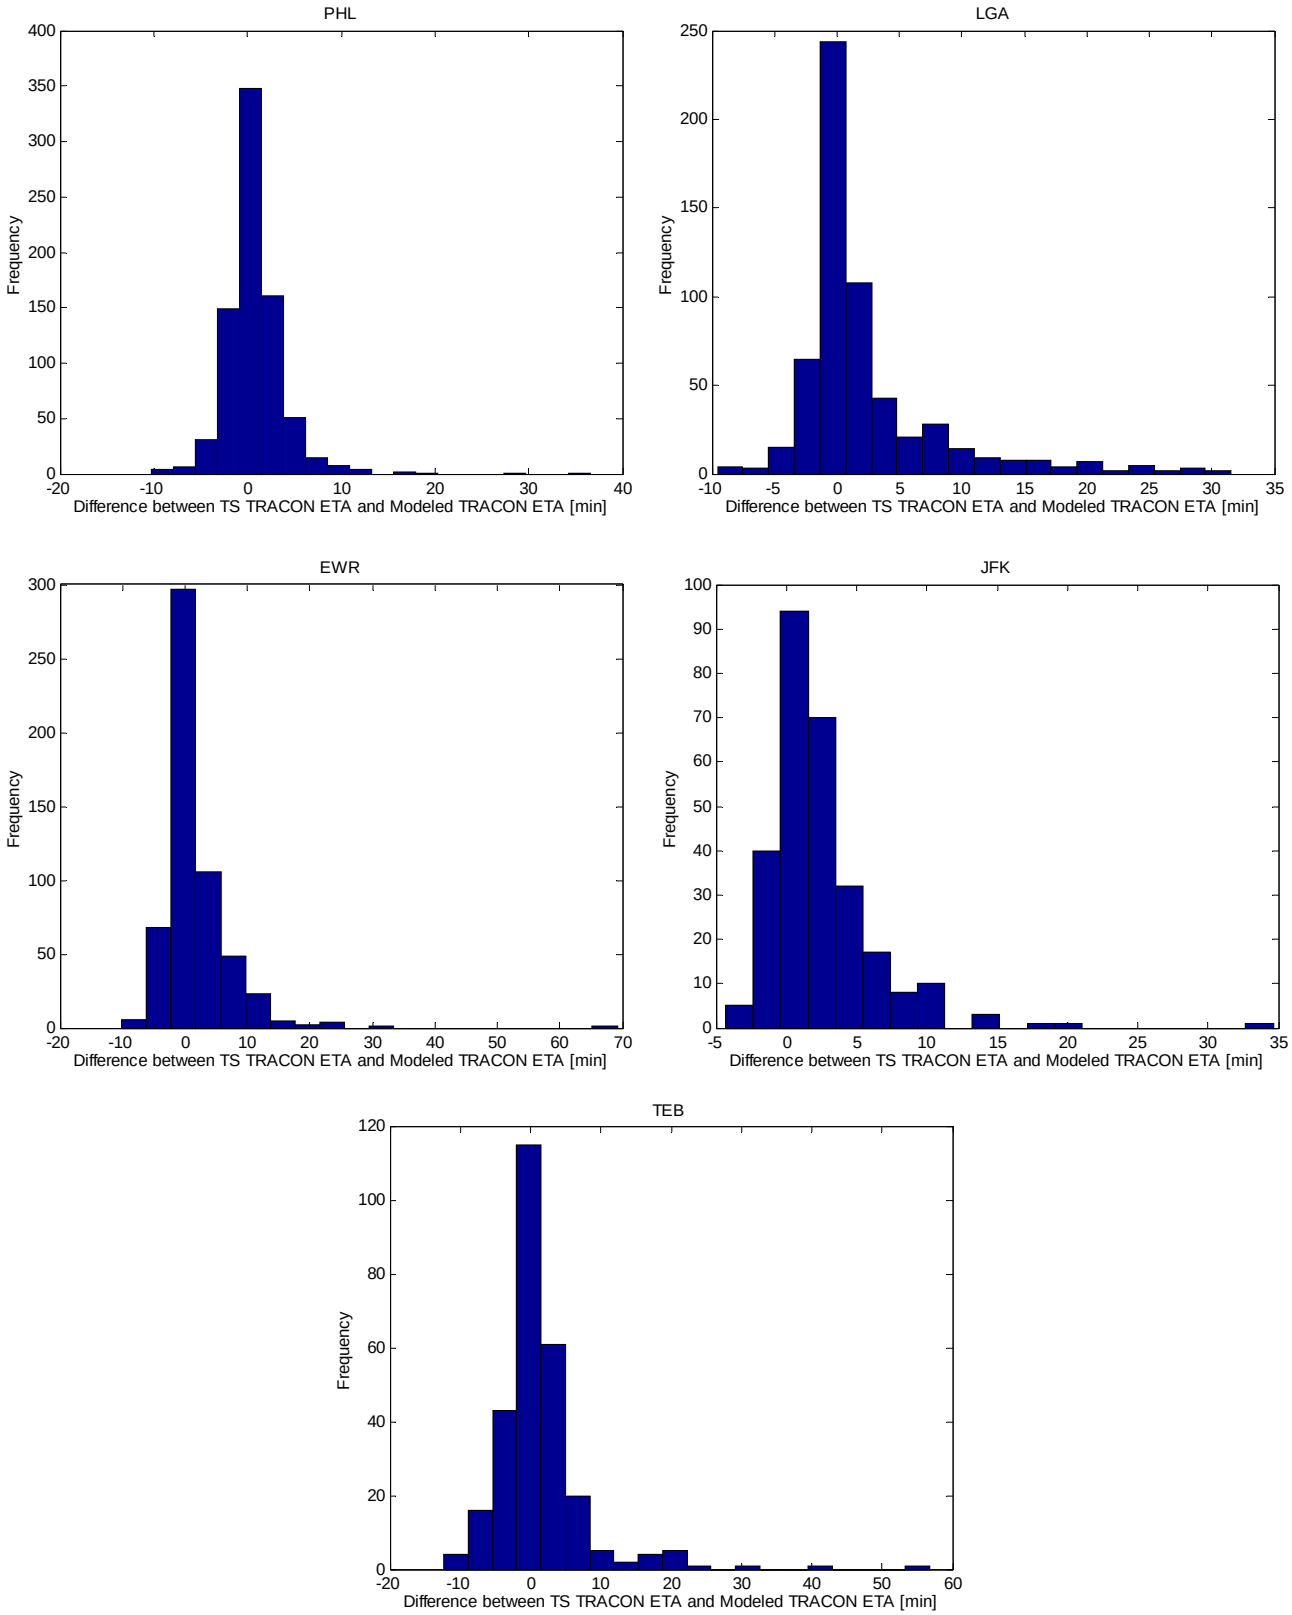

Figure 23. Histograms for presenting the differences between flight ETAs calculated using the CTAS TS, and calculated using the statistical modeled presented above.

According to the histograms in Figure 23 there are some cases where the difference between the TS and statistically modeled ETA is particularly high, such as nearly 70 minutes for one flight at EWR. These cases are likely to be errors in the flight plans, or in the TS parsing of the flight plans. However, the lower values in the tail, such as those at JFK and LGA, are not so high as to clearly be the result of errors, and may thus be caused by other effects. One such effect is short cuts from the scheduled flight plan. Figure 24 shows a plot of TS output tracks and corresponding host tracks (the tracks actually flown by the flights modeled by the TS) for the arrival flows into PHL. The instances where the flight paths actually flown deviate from the flight plan can easily be identified. This is particularly prevalent at the fix COFAX and HAR on the west arrival flow, and at the fixes CANNY and HEDGE on the west of the two south arrival flows, into PHL, as shown in Figure 24. The TS tracks show flights plans passing through COFAX, HAR, HEDGE and CANNY. However, the host tracks show that many flights did not in fact pass through these fixes, but took short cuts past them. It is possible that the flight plans were amended to exclude these fixes, but because the flight plans inputted into the TS were those at each flight's entry into the system, such flight plan amendments were not captured by the TS. The result is that ETAs calculated by the TS are later than actually flown. The statistical models were, however, developed from actual data, and thus account for any short cuts flown regularly. The statistically modeled ETAs are thus likely to be earlier in these cases than those modeled by the TS, as is the case in Table 2.

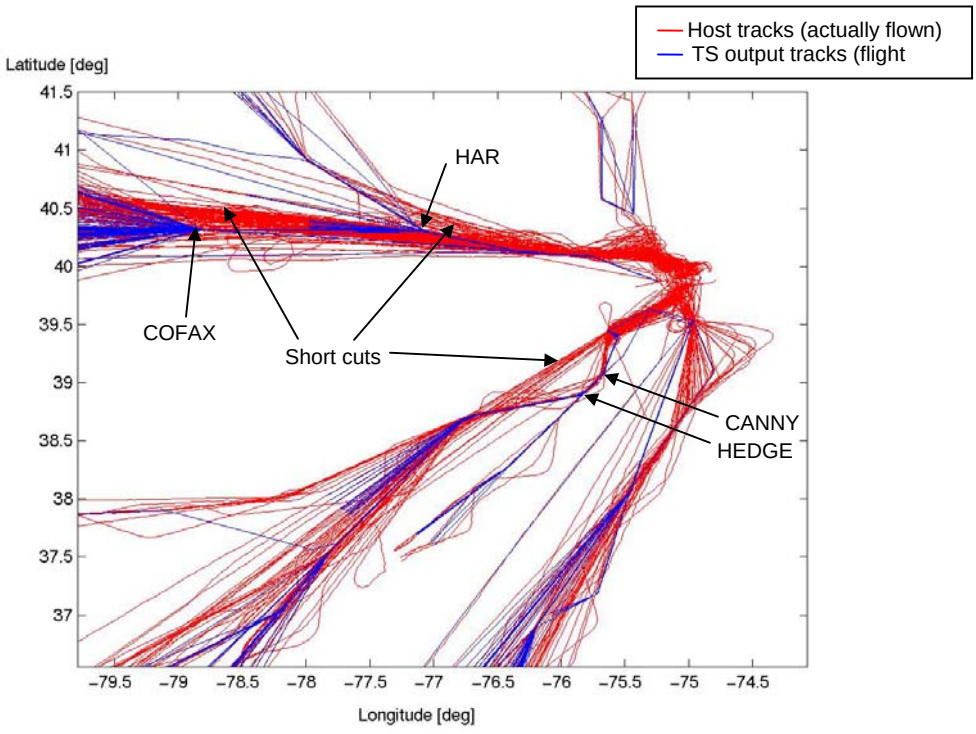

Figure 24. Host and TS output tracks arriving into PHL, showing short cuts flown relative to flight plans

Figure 25 shows the host tracks for arrival flows into BUNTS, the west arrival fix at PHL. The flights identified as unimpeded in the statistical modeling for the unimpeded

transition time between the boundary of ZNY and ZOB and BUNTS, and thus used to identify the unimpeded transition time, are highlighted. It is clear from these highlighted flights that they cover a range of paths between the meter fixes, including short cuts.

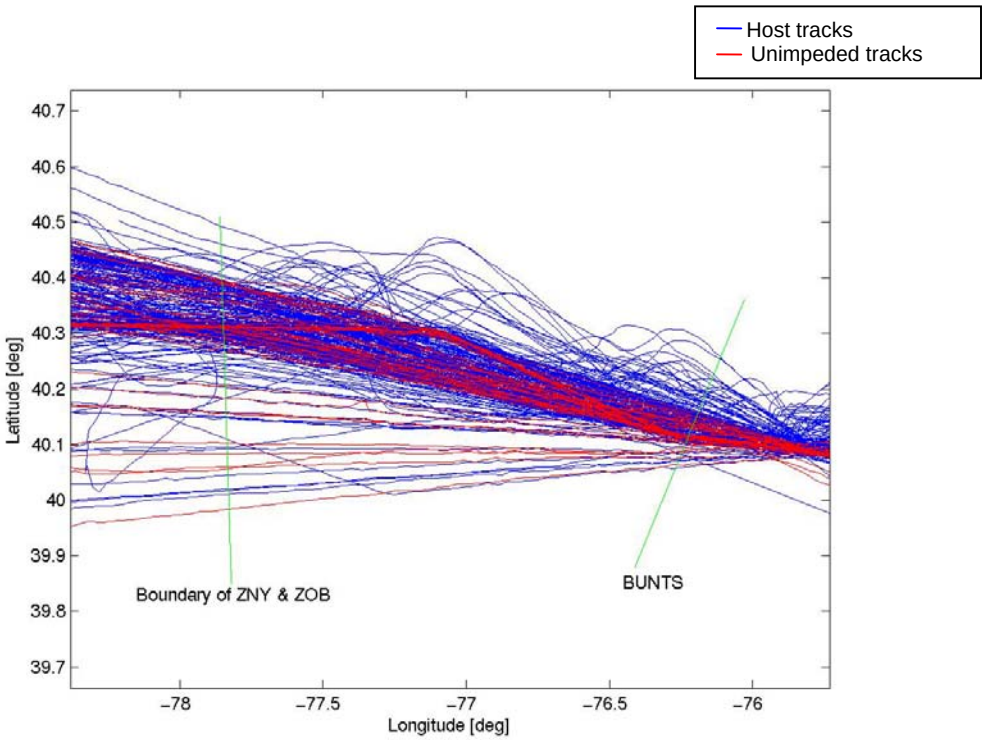

Figure 25. Host tracks for the west arrival flow into PHL from boundary of ZNY and ZOB to BUNTS, highlighting the flights identified as unimpeded in the statistical modeling of unimpeded transition times.

4. Identification and Modeling of Current Flow Management Procedures

When the traffic demand is expected to exceed the capacity of the system (whether determined by runway or TRACON acceptance rates) the air traffic managers apply a number of flow management procedures in order to avoid gridlock and excessive delays. The main flow management procedures concerning the PHL and N90 flows were identified through expert elicitation at the facilities visited. They are identified in order to understand the causes of the baseline behavior and limitations as well as to support the generation of a model of the current operations baseline. This model is described in this section and the parts that were implemented and simulated are indicated where applicable. Certain parts of the current operations that were identified through the site visits but were not modeled are described to provide insight and to support future research.

4.1. Miles in Trail and its propagation

Arrival flows are fed into N90 from ZNY, ZDC and ZBW. Arrival flows into N90 from ZDC and ZBW transition directly into N90, and do not pass through ZNY, with the exception of JFK arrivals from ZDC, which pass through ZNY before transitioning to N90. Arrivals from ZOB pass through ZNY before transitioning to N90. The PHL arrival flows are fed into PHL from ZNY and ZDC. When there are acceptance capacity constraints at the airports or in the TRACON restrictions are imposed on the inbound flow through Miles In Trail (MIT) at the arrival fixes. These MIT restrictions are propagated upstream from center to center when the delay required by a center is beyond its delay absorption capacity. This is illustrated schematically in Figure 26.

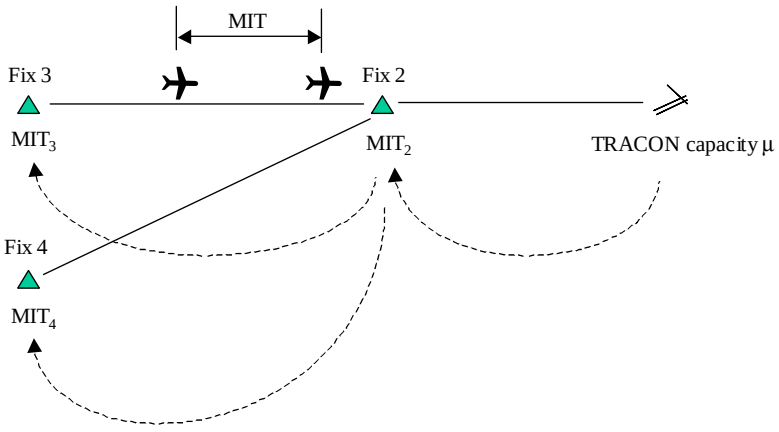

Figure 26. Propagation of Mile in Trail restrictions

ZNY is not able to absorb much delay and suffers from compression due to the speed and altitude reduction required on descent into N90. Consequently MIT restrictions

from N90 are normally passed directly to ZOB, ZBW and ZDC increased by 5 miles to account for compression. According to expert elicitation from TMCs at ZBW, ZOB and ZDC, restrictions are generally only propagated from ZBW to ZOB, from ZOB to ZAU, and from ZDC to ZTL, ZID and ZJX if the restrictions from N90 are greater than 15 MIT. Restrictions of 15 MIT or less are possible to accommodate by absorbing the delay within the centers airspace.

According to the interviews with TMCs the constraints within the ARTCCs that impact MIT propagation upstream are domestic upstream demand, internal departure demand, international arrival demand, wind, and the delay absorption capabilities of each sector. The restrictions propagated from ZBW to ZOB are primarily a function of the restrictions in place between N90 and ZBW, international demand from Europe, domestic demand from ZOB, and internal departure demand from Boston Logan (BOS), Manchester (MHT), and Providence (PVD) airports. The restrictions propagated from ZOB to ZAU are primarily a function of the restrictions in place between ZNY and ZOB, internal departures from Cleveland, Detroit and Pittsburgh, and demand from ZAU. The restrictions propagated from ZDC to ZTL, ZID and ZJX are primarily a function of the restriction in place between N90 and ZDC, internal departure demand from Washington Dulles (IAD), Washington National (DCA), Baltimore (BWI), and Philadelphia (PHL) airports and demand from ZTL, ZID and ZJX.

Offloading flights is used in conjunction with these MIT restrictions in certain cases. When restrictions from ZNY to ZOB are 20 MIT or more, TMCs at ZOB may offload some traffic through ZBW or ZDC, from which it can then enter N90. With offloading less MIT restrictions are passed to the upstream centers. This approach is only effective if the ZNY airspace is the constraint. When the N90 airspace is the constraint, the aircraft are simply delayed on ZBW/N90 boundary as opposed to the ZOB/ZNY boundary. Lack of coordination reduces the effectiveness of offloading.

When flows merge, restrictions are often passed back to more than one upstream flow. Lower MIT are generally imposed on the heavier flow to avoid starving the downstream resources, and because the inter spacing between aircraft from the lighter flow is larger needing a larger MIT to produce an effect on the flow rate. The accuracy with which this balancing is done is limited however by the low resolution of MIT restrictions (in increments of 5 MIT) and by the lack of knowledge of the downstream cause (for example, as indicated in an interview, an aircraft heading to a non-restricted runway does not need to be delayed, but not knowing the runway assignment does not allow favoring such an aircraft).

According to interviewed TMCs, MIT restrictions greater than 30 or 40 are rarely used. This is because such restrictions are difficult to apply effectively, and lead to unpredictability in the flow, and can lead to instability in the system. Consequently, holding, a Ground Stop, or a Ground Delay Program is implemented instead of such severe MIT restrictions. Typically, holding is applied first for an immediate effect, and then a Ground Stop is used as a temporary relief until a long-term Ground Delay Program is implemented.

A model of MIT and its propagation is shown in Figure 27. This model consists of two main components: A restriction generation model and a delay flow model. The

restriction generation model predicts the restrictions imposed by the facilities affecting arrival flows into New York and PHL by consideration of predicted AAR at the airports, and predicted arrival demand both at each airport as a whole, and on each airport's individual arrival fixes. The model's outputs include MIT to be applied on arrival fixes and MIT as propagated at facility boundaries. Other restrictions such as Ground Delay Program (GDP) and Ground Stop (GS) are modeled in terms of their interaction with MIT. The relationship between the inputs and the outputs is derived based on historical data and represents current procedures.

The restrictions are an input to the delay flow model, which takes these restrictions and converts them into separations to be applied between aircraft. These separations are determined by a statistical model based on historical data and represent current ATC behavior in meeting the assigned restrictions. Flight data is the arrival demand input to the delay flow model. Flights fly between different fixes along a flow network as described in Section 3.2. Each flight is assigned estimated times of arrival (ETAs) at fixes and at the TRACON boundary, by assuming unimpeded flight between meter fixes. The unimpeded flight time between each meter fix pair in the flow network is sampled from transition times of flights derived from historical data, as described in Section 3.2. The resulting ETAs are modified according to a statistical distribution of separations required by the restrictions, resulting in actual times of arrival (ATAs) at the TRACON boundary. These are converted to ATAs at the runway by the TRACON model, which is discussed in Section 4.2.

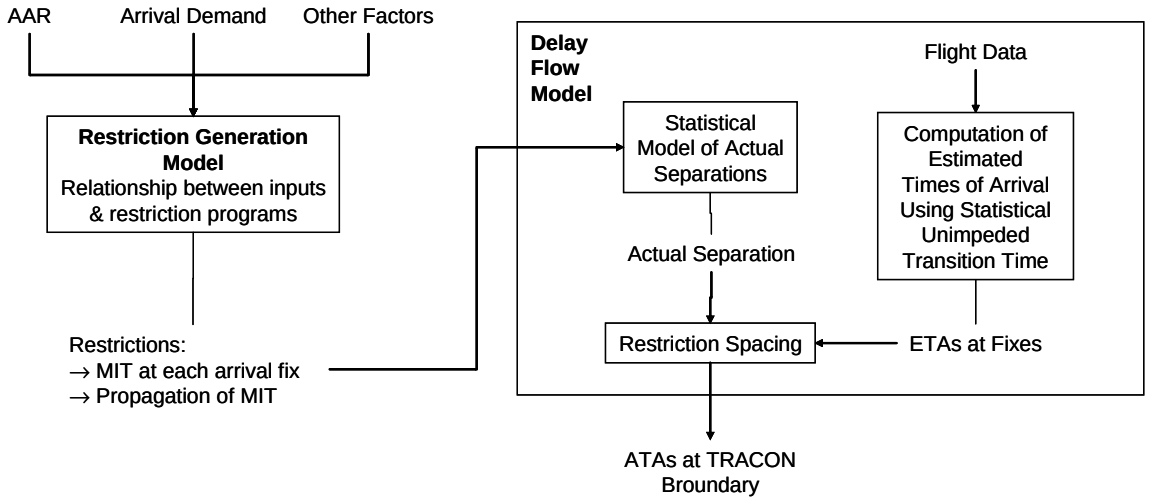

Figure 27. Current operations model diagram for MIT and its propagation

4.1.1. Restriction Generation Model

The restriction generation model determines MIT restrictions according to current procedures. Its underlying relationships were determined through analysis of one month of facility logs – the month of November 2003. Resulting restrictions meet reported AARs at airports, given demand levels on arrival fixes, and total airport demand.

4.1.1.1. Omitting periods with Ground Delay Program and Ground Stop

The occurrence of various restriction programs was identified from the logs and grouped by AAR and demand, available for the month of November from ASPM data. The purpose of this analysis was to identify how the various restriction programs were used with respect to one another.

Figure 28 and Figure 29 show how the restriction programs identified were used, as a function of AAR and demand/AAR respectively at PHL. The demand used in Figure 29 is ASPM demand, which includes both flights having landed in the time period (15 minutes in duration) and flights scheduled to have landed, but which have not yet landed. This measure of demand thus includes the effect of queuing. Restriction programs shown include MIT, GDP and GS. In each figure, a) shows non-normalized data in total duration of each program in minutes, to establish how large the sample set is for each bin shown in b), which shows the normalized data as percentages of duration. In both Figure 28 and Figure 29 one can see that when the airport was unconstrained – AAR was high and demand/AAR was low – 'no restrictions' was the dominant field. As the airport became more constrained – AAR decreased, and demand/AAR increased – first MIT was applied, followed gradually by more severe programs such as GDP and GS until at low AAR and high demand/AAR, MIT was not used on its own at all, but only in conjunction with these more sever programs.

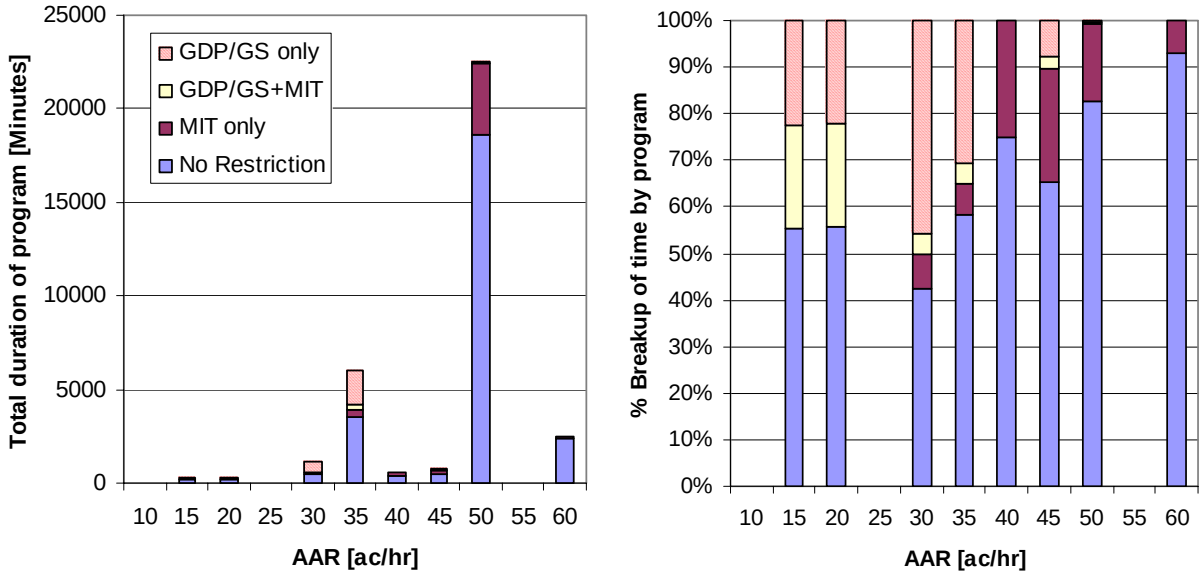

a) b)

Figure 28. Distribution of restriction programs versus AAR, at PHL, November 2003. a) Shows non-normalized data in total duration of program through the month, in minutes. b) Shows normalized data as percentages of duration.

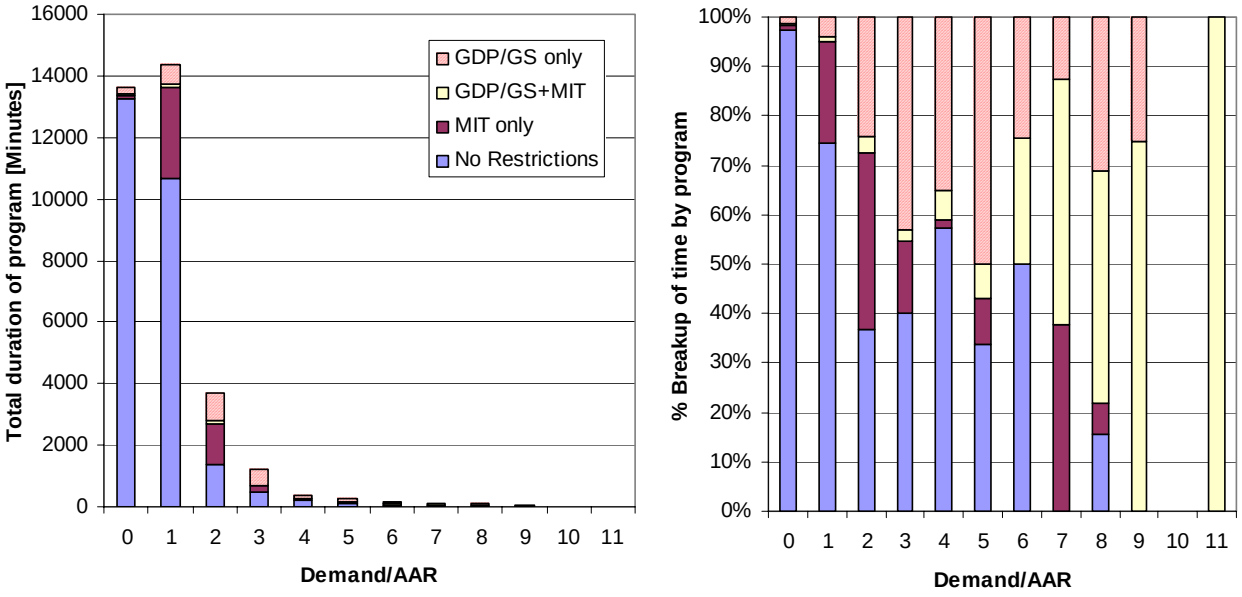

a) b)

Figure 29. Distribution of restriction programs versus Demand/AAR, at PHL, November 2003. a) Shows non-normalized data in total duration of program through the month, in minutes. b) Shows normalized data as percentages of duration.

In order to generate a model for MIT generation, periods of time with GDP or GS were omitted from the data since they would distort the causal relationship for applying MIT. These periods were highly constrained, yet had either MIT in place in conjunction with GDP/GS, or had no MIT at all. Furthermore, GDP and GS were not explicitly analyzed in this study.

The restriction model consists to two components: MIT prediction on the arrival fixes at the airports; and prediction of MIT propagation upstream from these arrival fixes.

4.1.1.2. MIT Assignment at Arrival Fixes

The process of MIT assignment at the arrival fixes was modeled by dividing it into two decisions. First the timing of the MIT restriction was determined, followed by the value of the MIT restriction.

The data available from the month of logs varied for each airport. Sample sizes are shown in Table 3 for each airport. These different sample sizes suggest that the models for different airports will vary in their accuracy. For PHL, for which the sample size is large, an accurate model was generated and the resulting model can be trusted with some confidence. However, at JFK, not enough data was available to establish a model that can be trusted with confidence. There was insufficient data to establish a relationship between different parameters and a model had to be borrowed from one of the other airports. The accuracy of the model for each airport reported later in this section will reflect the effect of the sample sizes.

Table 3. Number of cases of MIT identified across the arrival fixes and facility boundaries over the month of November 2003.

Airport	Number of Cases of MIT Identified
Philadelphia (PHL)	487
LaGuardia (LGA)	178
Newark (EWR)	118
Kennedy (JFK)	59
Teterboro (TEB)	86

Timing of MIT Restriction

In order to observe the correlation between demand and MIT restrictions, each was observed through the course of a day. Figure 30 to Figure 34 show the frequency of MIT applied on the arrival fixes at each airport as a function of time of day. Also shown is the demand (ASPM demand) at the airport as a function of time of day, averaged over the month. The analyses of both frequency and demand excluded days with GDP or GS at the airports in question.

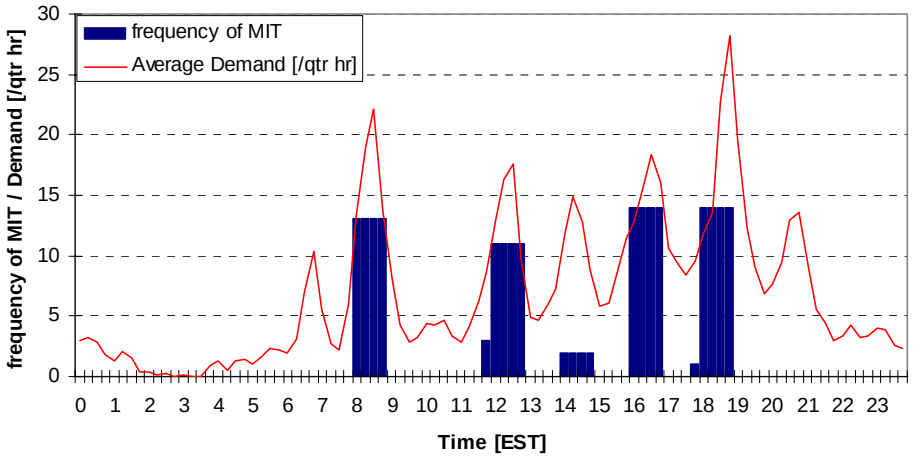

Figure 30. Frequency of MIT at PHL in the month of November, 2003, and average ASPM demand for each 15 minute time period through the day.

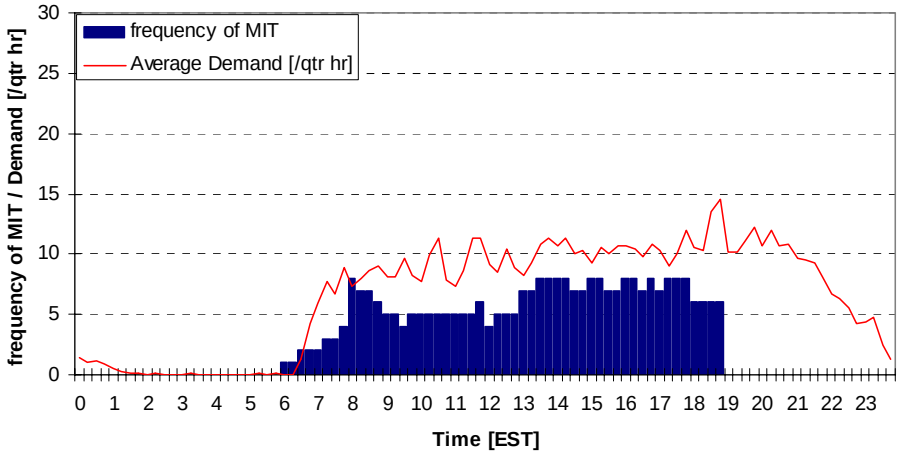

Figure 31. Frequency of MIT at LGA in the month of November, 2003, and average ASPM demand for each 15 minute time period through the day.

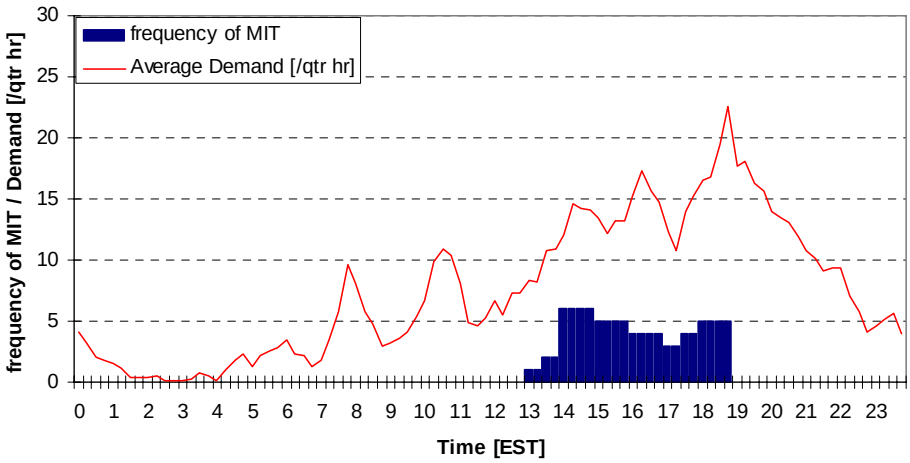

Figure 32. Frequency of MIT at EWR in the month of November, 2003, and average ASPM demand for each 15 minute time period through the day.

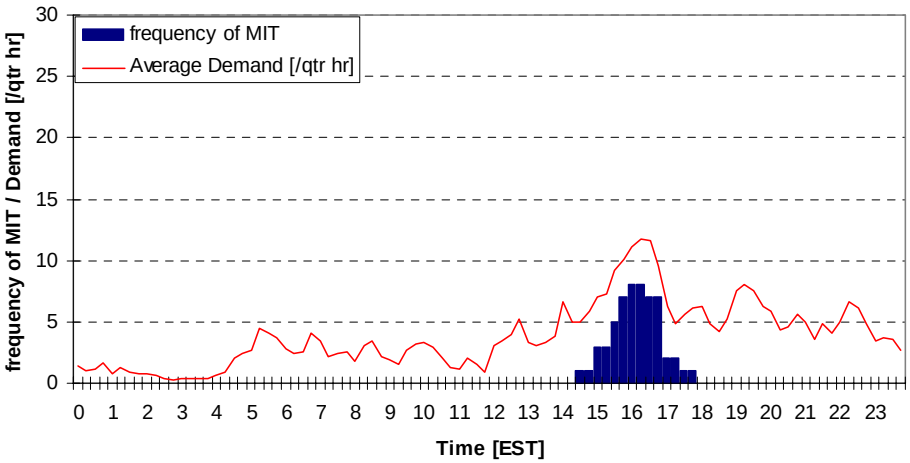

Figure 33. Frequency of MIT at JFK in the month of November, 2003, and average ASPM demand for each 15 minute time period through the day.

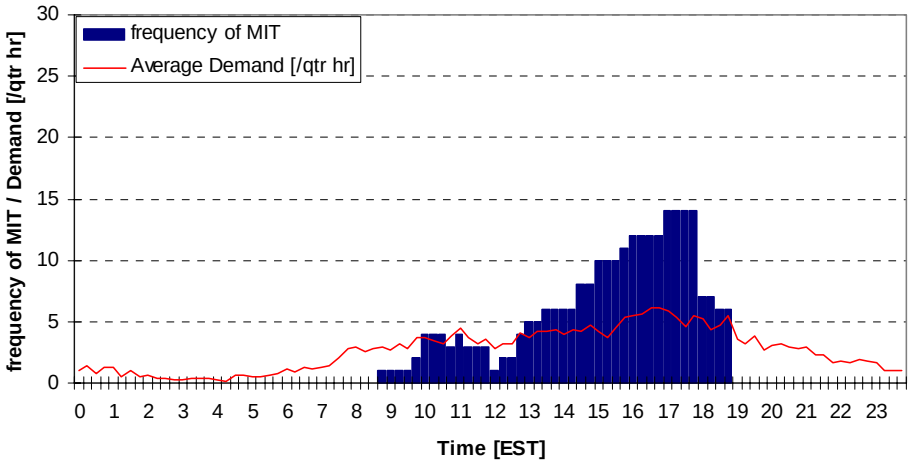

Figure 34. Frequency of MIT at TEB in the month of November, 2003, and average ASPM demand for each 15 minute time period through the day.

These figures show some correlation between MIT applied at arrival fixes, and demand at the airport. At PHL (Figure 30) there is clear correlation between MIT application and five of the arrival banks, which can be seen as spikes in demand. The timing of MIT application is also seen to be consistent through the month – the frequency is high and restricted to specific times of day, rather than being spread out over the whole day. At LGA (Figure 31), there is less indication of consistent times of MIT application, as frequency of MIT is spread out across the day with lower frequencies than were observed at PHL. The correlation with demand is still however apparent. The same can be said about EWR (Figure 32) and TEB (Figure 34). MIT application at JFK (Figure 33) all occurs during a short period in the afternoon, although it does not occur with high frequency. This period coincides with the afternoon arrival bank.

It is insightful to see the effect of GDP/GS on the presence of MIT. This effect is illustrated at PHL by comparing Figure 35 to Figure 30 above. Figure 35 was generated in the same way as Figure 30, but the periods of GDP and GS were not omitted.

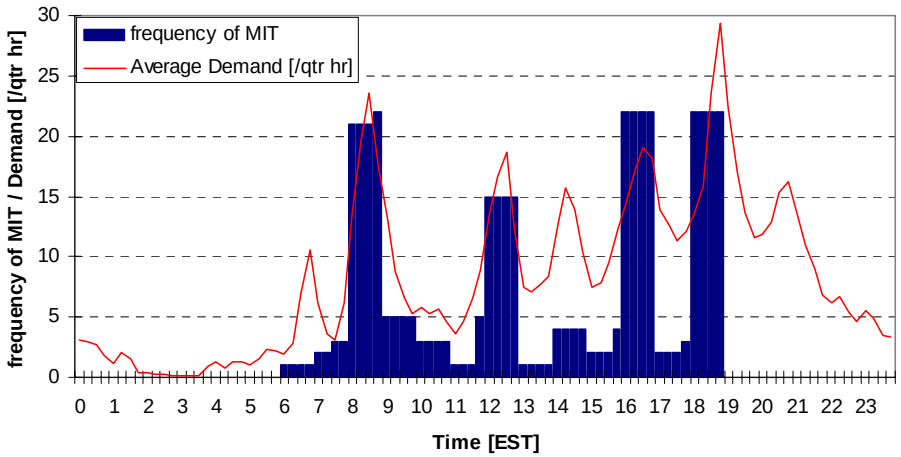

Figure 35. Frequency of MIT at PHL in the month of November, 2003 for each 15 minute time period through the day, for all days, including those with GDP & GS.

The same trends are apparent with the majority of MIT application falling on the arrival banks. This suggests that the inclusion of MIT during GDP/GS does not greatly affect the timing of MIT application.

In order to determine the parameter to characterize the timing of the application of MIT restrictions, a number of parameters were compared to the occurrence of MIT. Each parameter was compared to MIT application for each day separately, calculating a correlation coefficient for each day. The correlation coefficients were then averaged over all the days for which they were calculated. The correlation coefficients are shown in Table 4 for a number of parameters associated with demand and AAR. The parameters were calculated for each 15-minute period through the day using a one-hour moving window.

Table 4. Calibration coefficients between various parameters and the occurrence of MIT. Coefficients were calculated for each day and then averaged.

Airport	ASPM Demand	AAR	Scheduled Demand	Sch. Demand / AAR	Sch. Demand – AAR
PHL	0.679	0.077	0.649	0.646	0.642
LGA	0.465	0.105	0.415	0.410	0.408
EWR	0.528	-0.006	0.533	0.527	0.528
JFK	0.615	0.326	0.624	0.560	0.520
TEB	0.555	-0.228	-	-	-

There is no schedule at TEB, hence the absence of a Scheduled Demand and associated parameters reported for this airport.

Parameters which include demand correlate better than AAR. The difference between the correlation of ASPM demand and scheduled demand is small. This suggests that the schedule is as good a representation of when MIT were put in place as the ASPM demand, which includes the effect of queuing.

It shall be seen and discussed later that when looking at the value of MIT to be applied, scheduled demand no longer correlates as well as ASPM demand. Since it is likely that the facilities used the same parameter to decide on the timing as the value, the parameter chosen to characterize timing of MIT application was chosen as ASPM demand.

Having determined a characterizing parameter, the next step was to calculate a threshold in this parameter, which allowed definition of the start and end times of the restrictions to be applied. MIT is applied when the characterizing parameter is above the threshold and not applied when the characterizing parameter is below the threshold. A one hour moving window was used for the characterizing parameter. This window starts 15 minute before the period in which the restriction started (the start period) and

considers the ensuing hour, ending 30 minutes after the end of the start period. This smoothes the parameters and allows both duration and magnitude of the parameter to be represented, as it is a type of integral. The moving window used in shown in Figure 36 as window 2. Also shown are alternate moving windows 1 and 3 starting and ending at different times. The position of the moving window is important at PHL, where the spikes in demand are in the order of the length of the moving window. Window 2 skews the window slightly to the future, allowing the parameter to represent the forecast as well as the present.

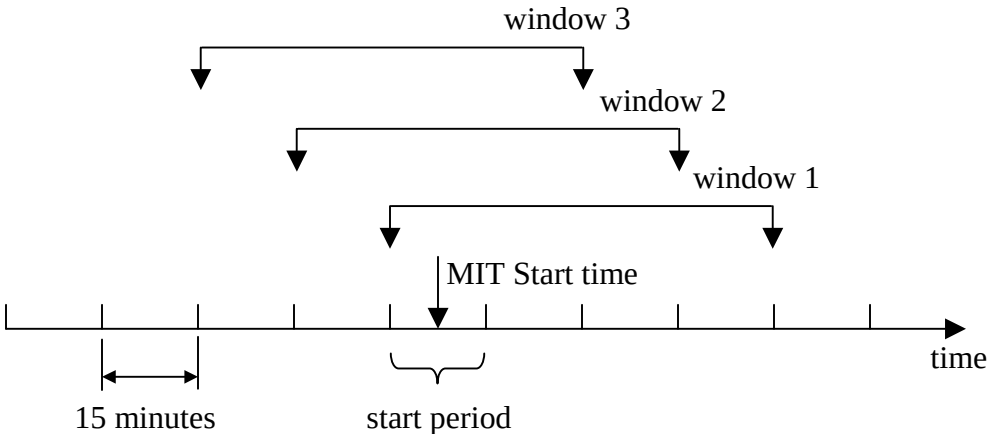

Figure 36. Schematic of moving windows used to calculate the smoothed value of ASPM demand as it was compared to the threshold.

The value of the threshold was varied and used to calibrate each model to the November data. In addition to the value of the threshold, two other characteristics of the model were varied as well. The first is the minimum time period for which the thresholds must be exceeded, for a restriction to be put in place. Since no restrictions shorter than an hour were observed – except some that were cancelled early – one hour was taken as the baseline. However, at PHL particularly, due to spikes in demand of short duration, the thresholds were exceeded for less than an hour, but still have MIT in place according to the log data. Since this minimum duration of demand exceeding the thresholds is likely to be different for different airports, it was varied from no minimum to 1 hour and used as a second characteristic of the model, to calibrate the models to the November data. (No minimum means that a restriction was put in place every time the threshold was exceeded even if it was only in one 15-minute period.)

Since no restrictions in the data were observed to be shorter than an hour, any restriction put in place for demand exceeding the threshold for less than an hour was extended to a full hour.

The second additional characteristic of the model varied was the minimum period of time between restrictions. If no limit is set on this time, a restriction could end only to start again in the next 15-minute time window. In the logs, restrictions were never observed to be separated by less than one hour, and so a minimum limit on this duration was needed. Again, this limit is likely to vary from airport to airport as at PHL for

example it is likely to be less than an hour, due to the banked structure, with banks sometimes only an hour apart – such as between the 16h00-17h00 bank and the 18h00-19h00 bank. This characteristic of the model was also varied between no minimum limit and 1 hour. (No minimum here means that restrictions were never merged even if one started in the 15-minute time period after another has ended.)

Therefore the threshold and the two time periods mentioned were all varied and modeled restrictions calculated for each combination of these characteristics of the model for the whole month. These restrictions were then compared to the actual restrictions in place in November 2003. This comparison allowed calibration of the model with the three characteristics selected for the most accurate model in the case of each airport. The metrics used are described below, and refer to Figure 37:

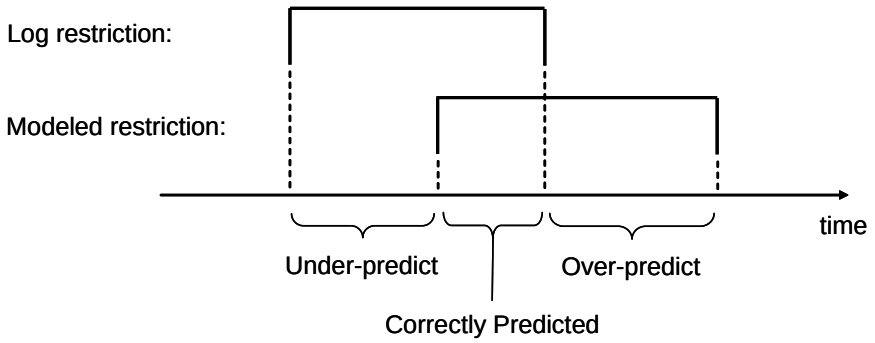

Figure 37. Comparison of modeled restriction with log restrictions.

- Overlap/(No overlap): The first metric used was the ratio of the total overlap between modeled restrictions and log restrictions (shown as "correctly predicted" in Figure 37); to the total time when the model was incorrect – either over predicting restrictions or under predicting restrictions (the sum of "over-predict" and "under-predict" as shown in Figure 37). This metric gave an indication of how accurate the model restricted time.
- (Over predict)/(Under predict): The second metric used was the ratio of restricted time over-predicted by the model to restricted time under predicted by the model ("over-predict" divided by and "under-predict" in Figure 37). This metric gave an indication of whether the model was over predicting or under predicting when it got restrictions incorrect.
- % of Restrictions correct: The third metric used wais the percentage of restrictions that are correctly predicted by the model. If there is any overlap between a modeled restriction and a log restriction, the modeled restriction is counted as correct. If there is no overlap, the modeled restriction is counted as incorrect.. Any log restriction that does overlap a modeled restriction is also counted as an incorrect prediction by the model. The modeled restriction in Figure 37 would be counted as correct.

The criteria for choosing the best characteristics included maximizing Overlap/(No overlap) while ensuring that (over predict)/(under predict) was close to unity. The % of restrictions correct was used as an indication the overall effectiveness of the model at predicting MIT.

The results of the calibration process are shown in Table 5 and Table 6. The accuracy of the model for the chosen characteristics is shown in Table 5 and is quoted as the three metrics. The threshold for each model, as well as the two minimum time periods varied for the calibration process are shown in Table 6.

Table 5. Calibrated accuracy of timing of MIT application on arrival fixes as determined from comparison with the logs for the month of November 2003.

Airport	Overlap/(No overlap)	(Over predict)/(Under predict)	% of Restrictions correct
PHL	1.50	1.12	71.8%
LGA	0.40	1.24	62.5%
EWR	0.41	1.89	64.7%
JFK	0.30	0.94	33.3%
TEB	0.22	1.45	35.5%

Table 6. Thresholds in demand for starting and ending MIT application and other parameters that allowed the best calibration of the models with November 2003 data.

Airport	Demand Threshold [ac/hr]	Minimum length of demand exceeding threshold	Minimum gap between restrictions
PHL	56	15 min	15 min
LGA	38	15 min	30 min
EWR	50	30 min	60 min
JFK	40	45 min	No minimum
TEB	26	45 min	60 min

The PHL, LGA and EWR models show better performance in terms of percentage of restrictions correct than the TEB and JFK models. This is expected from the sample sizes in each case (see Table 3). The overlap/(no overlap) metric is best at PHL, but better at LGA and EWR than at JFK & TEB. The banks at PHL made prediction of MIT easy, since demand was consistently high during these banks, and they match up well with observed restriction application (see Figure 30). At EWR and LGA, the MIT structure was not as regular, although the indication was that it still followed demand (see Figure 31and Figure 32). Demand however did not peak the way it did at PHL, and so accurate prediction of start and end times of restrictions was more difficult.

The thresholds calibrated to above are all within 14% of the medians of demand in the start period of MIT of the November data,

Value of MIT restriction

Having identified the timing of a restriction, the value of MIT by which the flights are to be spaced was to be specified at each arrival fix. The log data analysis allowed identification of a number of combinations of MIT applied on the arrival fixes at each airport. The number of times each of these combinations of MIT was observed is shown in Figure 38 to Figure 42 as a frequency, for all airports studied.

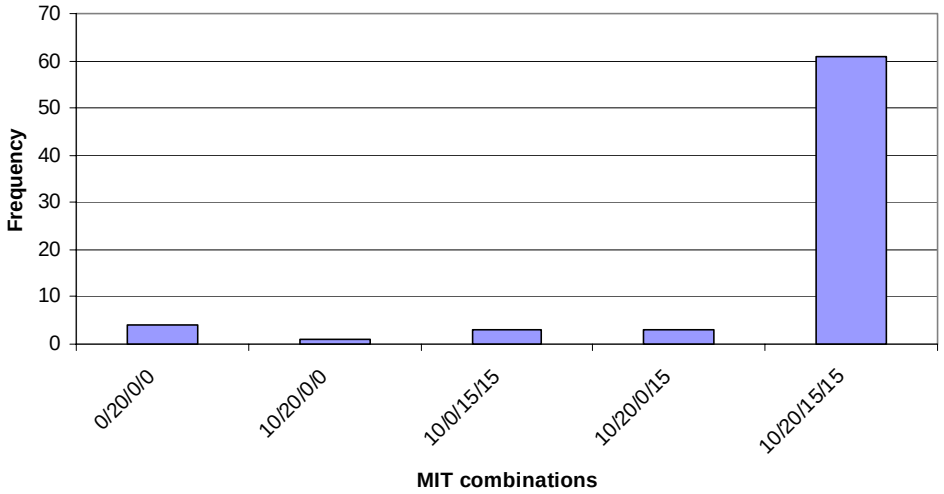

Figure 38. Frequency of different combinations of MIT on the arrival fixes at PHL, quoted as MIT at: BUNTS / SPUDS / VCN / TERRI.

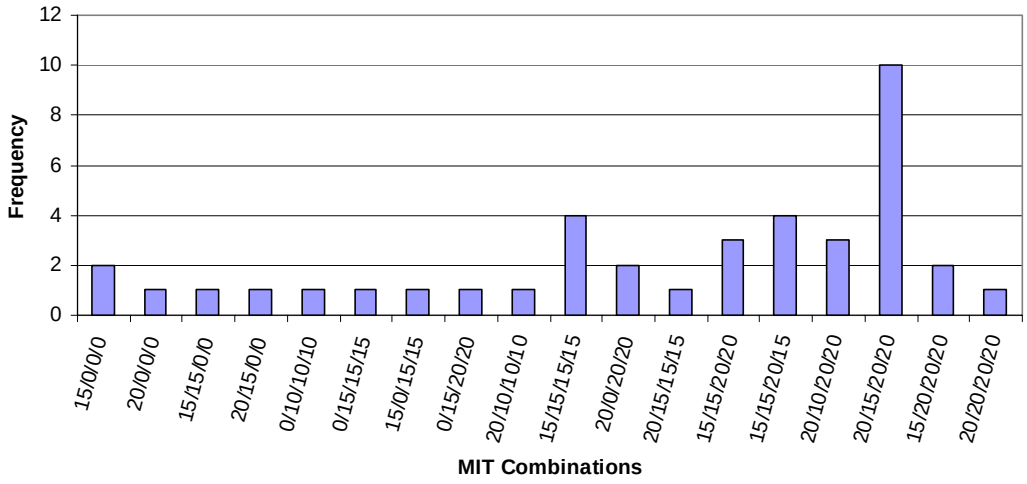

Figure 39. Frequency of different combinations of MIT on the arrival fixes at LGA, quoted as MIT at: LIZZI / RBV / NOBBI / VALRE.

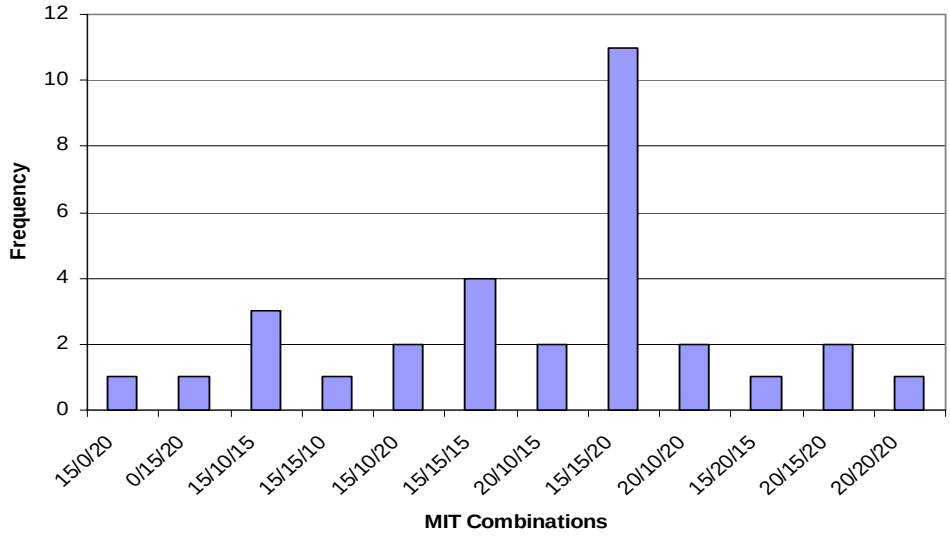

Figure 40. Frequency of different combinations of MIT on the arrival fixes at EWR, quoted as MIT at: PENNS / ARD / SHAFF.

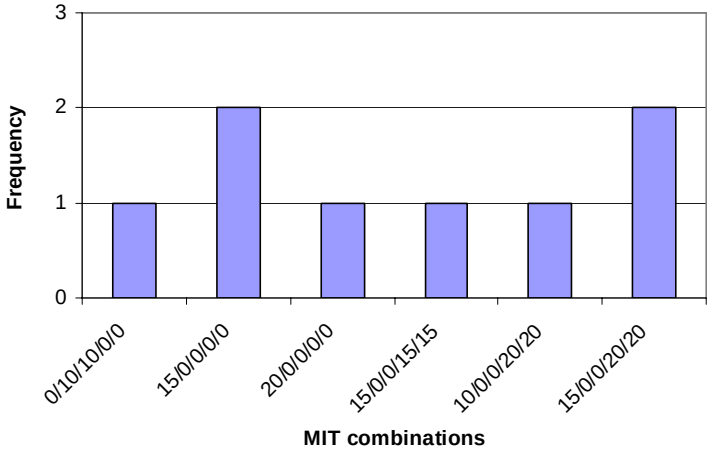

Figure 41. Frequency of different combinations of MIT on the arrival fixes at JFK, quoted as MIT at: LENDY / ZIGGI / CAMRN / CCC / LOVES.

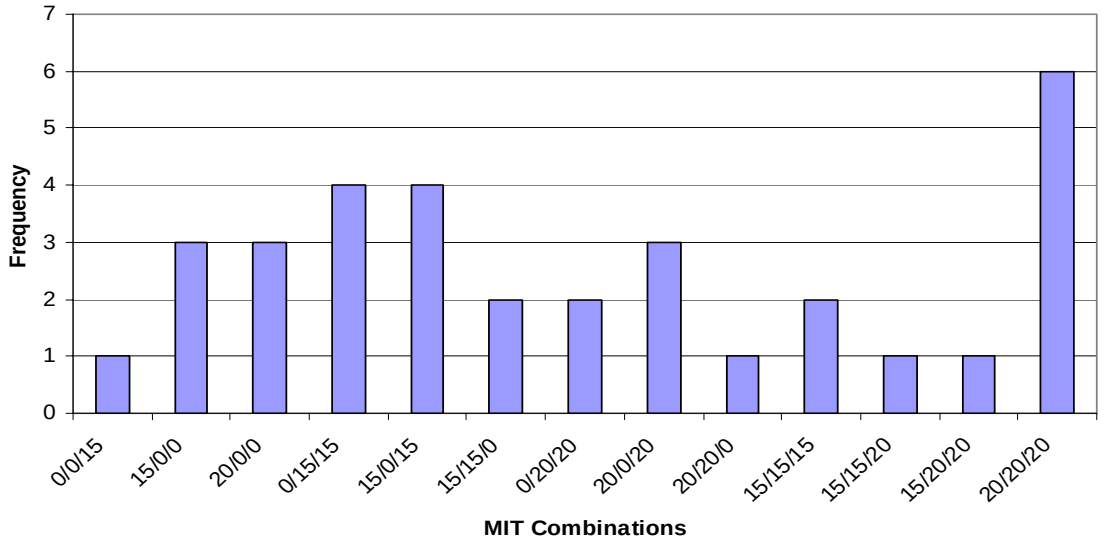

Figure 42. Frequency of different combinations of MIT on the arrival fixes at TEB, quoted as MIT at: MUGZY / MAZIE / LEMOR.

From these figures it can be seen that the number of combinations of MIT observed varies from airport to airport. At PHL and JFK for example, only 5 and 6 combinations were observed respectively. At JFK, this is due to a small sample size. At LGA, 20 combinations were observed. The sample size for each airport (as shown in Table 3) also affected how many occurrences were observed in each combination. At JFK for example, a maximum of 2 occurrences was identified for all 6 combinations, once GDP/GS time periods had been filtered out of the data.

The first step in the process to model the decision for determining the value of MIT was to establish a relationship between a characterizing parameter and the severity of MIT applied over the fixes during a given restriction. In order to remove the influence of outliers in the data, the calibration process to determine this relationship only used the data for the MIT combinations that occurred more frequently – any combination with more than two occurrences in the month of November was accepted for the calibration process while all others were discarded. In the event that the accepted MIT combinations did not cover the full range of severity of MIT, then particular combinations with one or two occurrences were accepted such that the full range of observed severity was represented.

At each airport, different combinations of MIT were identified as dominant. At PHL, one combination of MIT is clearly dominant, i.e.: 10 MIT on BUNTS on the PHL/ZNY boundary, 20 MIT on SPUDS, also on the PHL/ZNY boundary, 15 MIT on VCN on the PHL/ZDC boundary, and 15 MIT on TERRI, also on the PHL/ZDC boundary.

Having identified the dominant combinations, the MIT values across the arrival fixes were summed. This sum of MIT across the arrival fixes is a measure of the severity of the MIT restriction applied. It was to this sum of MIT that various characterizing parameters were correlated. This was done by calculating the average value of a number

of parameters incorporating ASPM demand, scheduled demand and AAR, over each restriction. These averages were then plotted against the sum of MIT across the arrival fixes for each restriction. The plot of the sum of MIT against ASPM demand at EWR is shown in Figure 43.

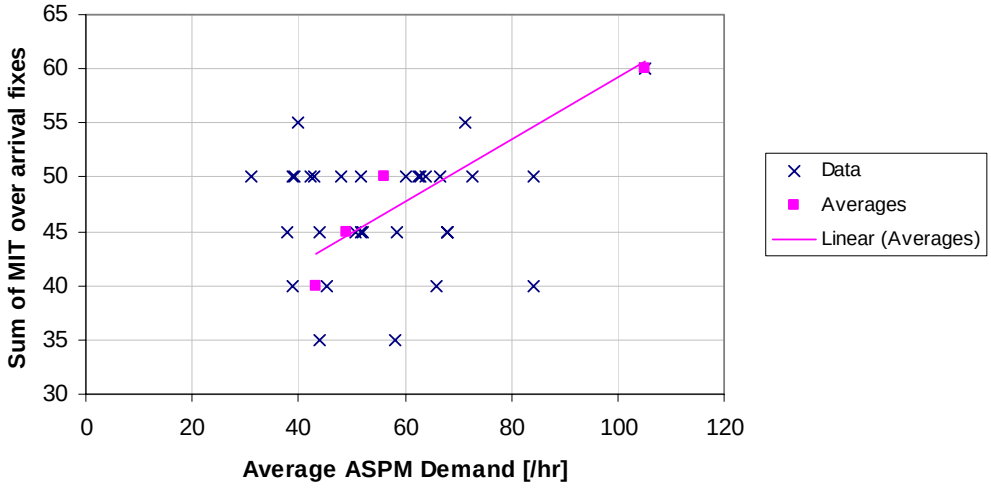

Figure 43. ASPM demand versus sum of MIT at EWR, with a straight line fitted to the averages of the ASPM demand.

At each level of the sum of MIT over the fixes, the parameters were again averaged and a line was fitted to the averages. The R^2 values of these curve fits were compared to choose the characterizing parameter. Table 7 shows the resulting R^2 values for a variety of parameters.

Table 7. R^2 values for various parameters plotted against sum of MIT values across the arrival fixes

Parameter	R^2 values				
	PHL	LGA	EWR	JFK	TEB
Average ASPM Demand	0.988	0.773	0.905	0.150	0.877
Maximum ASPM Demand	0.990	0.665	0.924	0.160	0.740
Average Scheduled Demand	0.128	0.226	0.139	0.079	-
Avg. Sched. Demand / AAR	0.340	0.010	0.404	0.075	-
Avg. Sched. Demand – AAR	0.405	0.028	0.532	0.096	-
Maximum Scheduled Demand	0.006	0.073	0.146	0.289	-
Max. Sched. Demand / AAR	0.442	0.144	0.070	0.256	-
AAR	0.264	0.603	0.852	0.063	0.048

The R^2 values in Table 7 show clearly that although ASPM demand and scheduled demand both correlated well with MIT when just considering the timing, when looking at the MIT value, scheduled demand no longer correlated well. ASPM demand, which reflects queuing effects, still correlated well. At this increased level of detail, the schedule was no longer a reasonable representation of the traffic on which the MIT decisions were based.

At TEB, no schedule is used, so no schedule data was available. At EWR, LGA and TEB average ASPM demand correlated best. Maximum demand correlated slightly better at PHL. At JFK, the correlation was poor for all parameters, but the sample size, as mentioned above, was small, so the low R^2 value was not necessarily meaningful. Based on these R^2 values, the parameter chosen to characterize MIT value was average ASPM demand.

Having determined the characterizing parameter, the relationship between this parameter and the sum of MIT values on the fixes was taken as the straight line fitted to the data (shown for EWR in Figure 43 above). For PHL however, despite the good correlation, it was also apparent that a single combination – 10/20/15/15 – was used on 86% of observed occasions (see Figure 38). The line fitted to the data is shown in Figure 44. The data indicates that a sum of MIT of 60 (which equals 10+20+15+15) is the maximum observed sum of MIT, and that, independent of demand, MIT is not only capped, but normally applied at this value. It is therefore reasonable to apply a sum of MIT of 60 independent of demand. Such a model would yield an accurate prediction of 2003, but would not allow for any higher MIT when demand increased in future years. In order to allow a higher MIT in extreme cases, it was decided to allow 20/20/20/20 to be applied – i.e. the most severer possible MIT combination. The demand above which this MIT would be applied is determined by extrapolating the fit line to a sum of MIT of 80. This yields a demand of approximately 80 ac/hr. So, at PHL, 10/20/15/15 is applied independent of demand, unless ASPM demand exceeds 80 ac/hr, in which case 20/20/20/20 is applied.

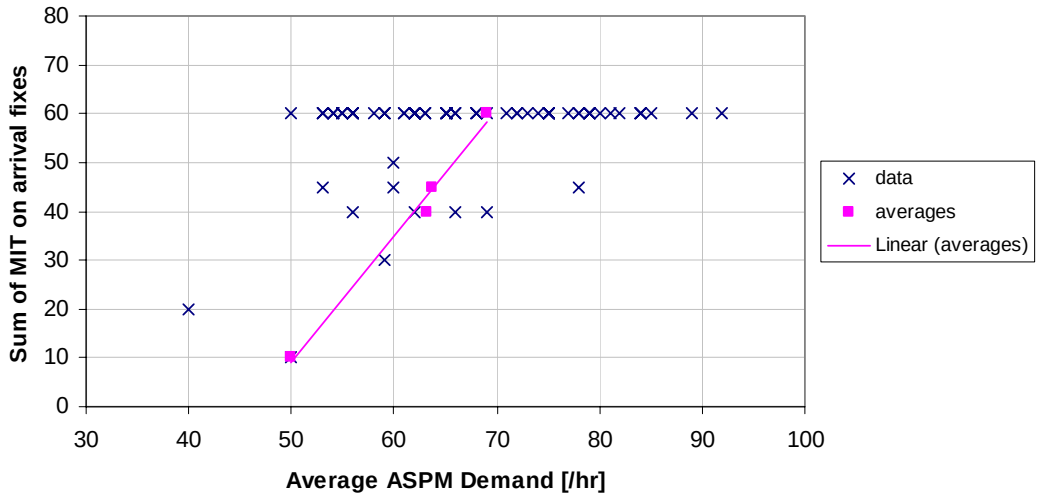

Figure 44. ASPM demand versus sum of MIT at PHL, with a straight line fitted to the averages of the data.

AT JFK, there was simply not sufficient data to build a reasonable model. Since sufficient data was available at other airports, it was decided to use a relationship developed at one of the other airports, to model JFK. EWR is most similar in terms of having an afternoon bank structure, and so the EWR relationship between ASPM demand and the sum of MIT on the arrival fixes was chosen to model JFK. The fact that JFK has 5 fixes whereas EWR only has 3 fixes is solved by the fact that MIT on CAMRN and ZIGGI were always observed to be the same; and MIT on LOVES and CCC were also observed to always be the same. So the 5 arrival fixes at JFK are narrowed down to 3 – LENDY, ZIGGI/CAMRN and CCC/LOVES.

Once the sum of MIT on the arrival fixes had been determined, the specific MIT value on each fix was to be determined. There was an indication from the data that lower MIT was placed on fixes with high demand, while higher MIT was placed on the fixes with low demand. For example, over the month, 38.5% of demand at PHL arrived through BUNTS, 7.4% arrived through SPUDS, 26.7% arrived through TERRI and 27.5% through VCN. The commonly used MIT combination at PHL was 10 MIT on BUNTS – the fix with the highest demand – 20 MIT on SPUDS - the fix wit the lowest demand – and 15 MIT on each of TERRI and VCN – the two fixes with medium demand. This rule was found to accurately describe 85% of observations of MIT at EWR, 53% of observations at LGA, 89% of observations at TEB, 20% of observations at JFK, and >86% of observations at PHL. This rule was also confirmed by the procedure elicited during the site visits, and was therefore chosen to model specific MIT on the fixes.

No MIT was observed to exceed 20 miles on the arrival fixes. It was decided to cap the MIT on the arrival fixes at this value, since even with expected increases in demand in future years, it was not expected that higher MIT would be used on the arrival fixes, because it would result in excessive MIT propagated upstream.

For an airport with a given number of arrival fixes, each sum of MIT has a number of possible combinations of MIT over that number of arrival fixes. Not all of these combination were however observed in the data. However, just because a combination was not observed, does not mean that it could not happen, and so the model was written to apply MIT according to the rule established with the relative demand on the fixes, rather than according to the observed combinations of MIT only. An example follows:

For a sum of MIT = 40 at EWR, where there are 3 arrival fixes, the following alternative combinations of MIT on the three fixes are possible:

1. two 15s, one 10
2. two 20s, one 0
3. one 20, two 10s

Alternative 1 was observed at EWR. Alternatives 2 and 3 were not however observed. The model still allowed alternatives 2 and 3 to be applied, by following the logic described following. This logic is based on observations of MIT application in the data, and the rule mentioned above:

When the sum of MIT across the arrival fixes is 40 miles:

1. If the fix with the lowest demand has less than 10% of demand, apply 20 MIT (the maximum MIT) on this fix, and 10 MIT on the other two fixes.
2. If the fix with highest demand has more than 90% of demand, apply nothing on this fix (the minimum MIT), and 20 MIT on each of the other two.
3. Otherwise, apply 15 MIT to the two fixes with lowest demand and 10 MIT to the third fix.

In this way, all MIT combinations were made possible, with decisions being made based on ratio of demand on each fix relative to total demand on the airport.

4.1.1.3. MIT Upstream Propagation

The propagation of MIT was modeled as a set of rules of the following form: if MIT at arrival fix XYZ is '*a*', then the propagated MIT to the upstream facility is '*b*', starting '*m*' minutes before the arrival fix restriction starts and ending '*n*' minutes before the arrival fix restriction ends. The propagation scenarios to be modeled were identified by considering the flow networks identified for each airport along with actual MIT restriction propagations observed in the month of November. Table 8 and Table 9 show the facility boundaries considered as a result of this identification analysis.

Table 8. Facility boundaries identified for propagation of restrictions at PHL & LGA

Airport	Arrival fix / facility	Facility boundaries for restriction propagation	
		1st propagation	2nd propagation
PHL	BUNTS / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	SPUDS / ZNY	ZNY / ZBW	
		ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	VCN / ZNY	ZNY / ZBW	
	VCN / ZDC	No Propagation	
	TERRI / ZDC	No Propagation	
LGA	LIZZI / ZNY	ZNY / ZOB	ZOB / ZID
	VALRE / ZBW	No Propagation	
	NOBBI / ZBW	No Propagation	
	RBV / ZDC	ZDC / ZTL	
		ZDC / ZJX	
		ZDC / ZID	

Table 9. Facility boundaries identified for propagation of restrictions at EWR, JFK and TEB.

Airport	Arrival fix / facility	Facility boundaries for restriction propagation	
		1st propagation	2nd propagation
EWR	PENNS / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	SHAFF / ZBW	ZBW / ZOB	
	SHAFF / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	ARD / ZDC	ZDC / ZTL	
		ZDC / ZJX	
		ZDC / ZID	
JFK	LENDY / ZBW	No Propagation	
	LENDY / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	CCC / ZBW	No Propagation	
	CCC / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	LOVES / ZBW	No Propagation	
	ZIGGI / ZDC	No Propagation	
	ZIGGI / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	CAMRN / ZDC	No Propagation	
	CAMRN / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
TEB	MUGZY / ZNY	ZNY / ZOB	ZOB / ZID
			ZOB / ZAU
	LEMOR / ZNY	No Propagation	
	LEMOR / ZBW	No Propagation	
	MAZIE / ZDC	ZDC / ZTL	
		ZDC / ZJX	
		ZDC / ZID	

The propagation rules were derived from a statistical analysis of observed occurrences of propagation in the November data. In some cases large sample sizes were available, such as at PHL, where there were 84 cases of 10 MIT on BUNTS propagated

to various values of MIT on the ZNY/ZOB border. Other scenarios had much smaller sample sizes.

Analysis of the data revealed that in most cases there was clear dominance of a particular MIT value propagated for each MIT value on the downstream fix or facility boundary. For example, 10 MIT on BUNTS was propagated to 20 MIT on the ZNY/ZOB boundary in 82% of the 84 occurrences of this scenario. The remaining 18% of occurrences were distributed between various other MIT values including no propagation at all. The means chosen to model propagation was to identify this dominant propagation rule, and apply it to all occurrences of the scenario in question. In this example therefore, the model would always propagate 10 MIT at BUNTS as 20 MIT on the ZNY/ZOB boundary.

In some scenarios, a dominant propagation rule was not as easily identified, as multiple options each had more than 30% of observed occurrences in the data. In such cases, all highly probable propagations were considered. The decision of which propagation to choose in the model was made randomly, based on the likelihood of each. For example, there were 76 occurrences of 20 MIT on the ZNY/ZOB facility boundary for flows through BUNTS to PHL. 42% of these were propagated to 30 MIT on the ZOB/ZID boundary. 50% were not propagated at all, and the remaining 8% were propagated to other values. In the model therefore, both a propagation of 30 MIT and no propagation were identified as dominant alternatives. The propagation applied was chosen randomly, with a 42% chance of propagating as 30 MIT and a 58% chance (50% + remaining 8%) of no propagation at all.

An example of the rules applied is shown in Table 10 for flows through BUNTS arriving into PHL.

Table 10. The propagation rules for MIT at BUNTS, propagated to the ZNY/ZOB facility boundary and the sample size on which each rule was based.

MIT at BUNTS	MIT Propagated	% use by model	% representation of sample set	Sample size
	to ZNY/ZOB			
10	20	100%	82%	84
	Other values	Not used	18%	
15	30 @ PSB, 20 @ J152*	100%	100%	1
20	40	100%	100%	1
	to ZOB/ZID			
10	30	100%	100%	1
20	30	42%	42%	76
	No Propagation	58%	50%	
	Other values	Not used	8%	
30	30	100%	-	No Data
40	30	100%	100%	1
	to ZOB/ZAU			
10	30	100%	100%	1
20	30	27%	27%	56
	No Propagation	73%	68%	
	Other values	Not used	5%	
30	30	100%	-	No Data
40	30	100%	100%	1

There is no data for the scenario of 30 MIT on the ZNY/ZOB boundary. The propagation value chosen was therefore based on trends in rules for other MIT values on this facility boundary.

Timing of the propagated MIT was similarly based on a statistical analysis of data from the month of November 2003. The upstream restriction usually preceded the downstream restriction, allowing the same flights to be metered at the downstream facility as at the upstream facility. The difference between upstream and downstream start times was calculated for all observed MIT propagations. The time difference with highest frequency on each facility boundary, fix and destination airport combination was chosen to be used in the model. Figure 45 shows the frequency of the difference in start time for restrictions propagated from BUNTS to ZNY/ZOB for flows arriving into PHL.

* PSB is a fix on the PHL arrival flow on the ZOB/ZNY boundary. J152 is a jet route that also falls on the PHL arrival flow, but is adjacent to the PSB flow. The two flows merge before reaching BUNTS.

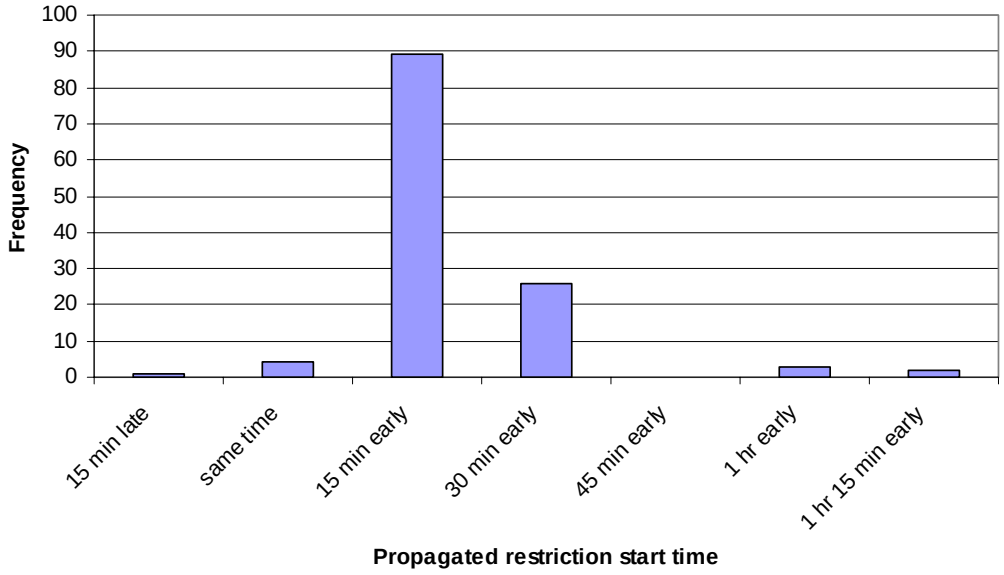

Figure 45. Frequency of occurrence of the difference between propagated restriction start time and arrival fix start time at BUNTS, for flows arriving into PHL.

In the above case, the propagated restriction was chosen to start 15 minutes earlier than the arrival fix restriction. The same analysis was performed for the end time. All results for flows arriving into PHL are shown in Table 11.

Table 11. The timing rules for propagation of MIT for arrival flows into PHL.

Fix	Arrival Fix		One facility out		Two facilities out	
	Start Time	**End Time**	**Start Time**	**End Time**	**Start Time**	**End Time**
BUNTS	*x*	*y*	*x* - 15min	*y* - 15min	*x* - 60min	*y* - 60min
SPUDS	*x*	*y*	*x* - 15min	*y* - 15min	*x* - 60min	*y* - 60min
VCN	*x*	*y*	*x* - 15min	*y* - 15min	No further propagation	
TERRI	*x*	*y*	No propagation			

In addition to MIT propagated from the arrival fix, MIT that was not propagated or associated with any MIT on any other fix or facility boundary, was also observed on facility boundaries,. These MIT were presumed to be in response to local constraints, such as sector overload. These MIT were not modeled in this baseline model and were excluded from the benefits analysis as described in Section 6.

4.1.1.4. Model Performance

The model was run at each airport for the selection of days in November 2003 for which contiguous host data was available, omitting days with GDP/GS. Host data was needed for demand on each of the arrival fixes. The fraction of these days restricted by the model was then compared to the fraction of the same period of time restricted in the logs. The results are shown in Table 12. Also shown is the number of days for which each fraction is calculated.

Table 12. Comparison of fraction of time in November 2003 restricted by the logs and by the restriction generation model respectively

Airport	Log restrictions	Model Restrictions	# days
PHL	13.7%	11.7%	10
LGA	12.8%	22.2%	11
EWR	6.3%	9.9%	4
JFK	1.8%	1.9%	14
TEB	11.7%	7.2%	15

PHL and TEB are under predicted by the model, whereas EWR and JFK are over predicted. LGA is also over predicted, but to a greater extent. This is due to relatively fewer restrictions in place in the days analyzed than is typical when the whole month is observed. The model is based on the whole month's data and so predicts more restrictions. The metrics used for calibration, as shown in Table 5, indicate that over the whole month, LGA over predicts by just 24% (over predict/under predict = 1.24.)

The use of the restriction generation model for extrapolation to future years is shown in Section 10.1.

4.1.2. Delay Flow Model

Arrival flows into an airport are delayed upstream of the TRACON by spacing aircraft according to MIT restrictions, and within the TRACON by delaying aircraft to meet runway capacity limitations. The delay flow model, a shown in Figure 27, is thus separated in two components, a MIT spacing model and a TRACON delay model. Each of these models, and their calibration against actual delays, are presented below.

MIT Spacing Model

Delay under baseline operations is calculated by applying the MIT identified by the MIT restriction generation model described in Section 4.1.1. The identified MIT are applied to the simulated arrival demand into each of the five airports studied, described in Section 3.2. Because MIT restrictions are applied on facility boundaries, the MIT restrictions calculated in the MIT restriction model are only applied at the meter fixes (at tier boundaries) shown in the arrival flow network in Figure 18 and Appendix A that coincide most closely with the facility boundaries. MIT are applied to the demand at these facility boundaries according to a delay flow model. This model is illustrated in Figure 46 below.

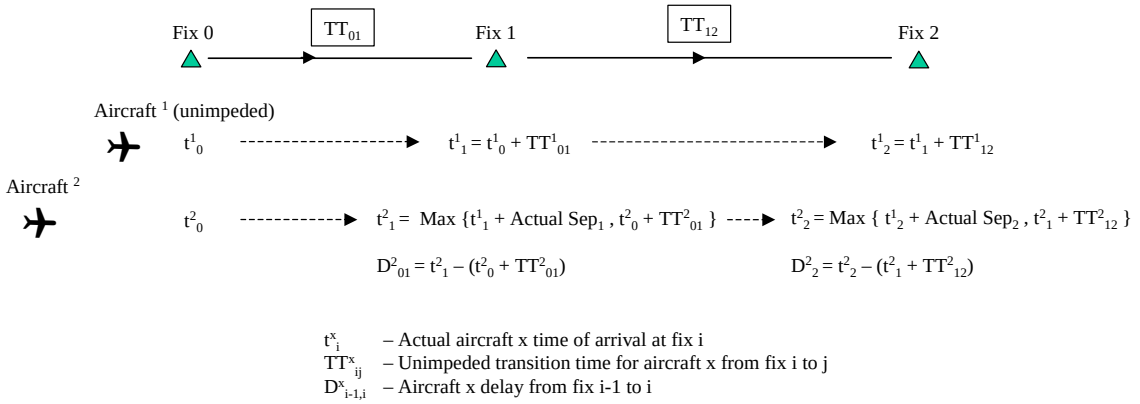

Figure 46. Calculation of meter fix crossing times and delay according to distance based restrictions (superscript = aircraft index, subscript = fix index).

The first aircraft in a demand bank is flown through each meter fix according to unimpeded transition times from meter fix to meter fix, as it should not be impeded by any aircraft in front of it, regardless of restrictions. Meter fix crossing times for all trailing aircraft are modeled relative to the first, unimpeded aircraft, according to the restrictions in place and the time each aircraft under question takes to fly between the fixes. This is done by calculating the aircraft arrival time at the meter fix were it to be separated from the leading aircraft in the queue by the specified MIT, and calculating the time of arrival at the meter fix if the aircraft flew unimpeded. The actual time of arrival at the fix is the greater of these two times. This is presented in equation (4) below:

$$t^{i}{}_{j} = \max\left\{ t^{i-1}{}_{j} + \frac{MIT}{v^{i}{}_{j}}, t^{i}{}_{j-1} + TT^{i}{}_{j-1,j} \right\} \quad (4)$$

where t^i_j represents the arrival time for aircraft *i* at meter fix *j*, and v^i_j represents the speed of aircraft *i* at meter fix *j*. $TT^i_{j,k}$ represents the transition time for aircraft *i* from meter fix *j* to meter fix *k*.

The resulting time of arrival is used for calculations of fix crossing times at the next fix downstream, and for the next aircraft in the queue, as illustrated in Figure 46.

The delay $D^i_{j-1,j}$ that must be absorbed by each aircraft, between each pair of fixes *j-1* and *j*, in order for the aircraft to meet the MIT requirements, was calculated as the difference between the modeled time of arrival at a fix, and the unimpeded time of arrival at the same fix. In the case of an unimpeded flight this delay would be zero. Equation (5) for this delay is as follows:

$$D^{i}{}_{j,j-1} = t^{i}{}_{j} - \left(t^{i}{}_{j-1} + TT^{i}{}_{j,j-1} \right) \quad (5)$$

where t^i_j and $TT^i_{j,k}$ are as defined above, and $D^i_{j,k}$ represents the delay incurred by aircraft *i* between meter fix *j* and *k*.

MIT Spacing Model Calibration

The MIT spacing model was calibrated by adjusting the error in meeting the specified MIT restrictions so as to equate the distributions of actual and modeled delay at the TRACON boundary as closely as possible. According to interviews with TMCs on the site visits actual separations are generally below the specified MIT restriction by between 0 and 2 Miles, although the numbers vary by controller. The mean error in meeting the specified MIT restrictions was adjusted so as to reduce the mean difference between the actual and modeled delay at the TRACON boundary.

Because the mean error is expected to increase as MIT restrictions increase in magnitude, the mean error was specified as a percentage of the specified MIT. This percentage was then adjusted to calibrate the delay spacing model against actual delay. The standard deviation in meeting the specified restriction was specified as a fixed number, of 2 nm, because it is not expected to increase as clearly with increasing MIT restrictions. The model could be calibrated further by adjusting this parameter. The intention was to adjust the standard deviation of the error in meeting the specified MIT restrictions, to reduce the difference between the standard deviations of the actual and modeled delay distributions at the TRACON boundary. This, however, has a second order impact on the results, and was not completed due to time constraints.

The results of the MIT spacing model calibration are presented in Table 13. In Table 13 $\mu_{\varepsilon\text{-MIT}}$ refers to the mean of the error in applying the specified MIT error, as a percentage of the specified MIT, that calibrated best. $\sigma_{\varepsilon\text{-MIT}}$ refers to the standard deviation of the error in applying the specified MIT error, in nautical miles. D_{act} refers to the average actual delay at the TRACON boundary, calculated according to actual flight arrival times at the TRACON boundary, relative to calculated ETAs at the TRACON boundary. $\mu_{\varepsilon\text{-D}}$ refers to the mean error between the modeled delay and the actual delay D_{act}. Similarly $\sigma_{\varepsilon\text{-D}}$ refers to the standard deviation of the error between the modeled

delay and the actual delay D_{act}. The ratio of the mean error $\mu_{\varepsilon\text{-}D}$ and the actual delay D_{act} is also presented, as a percentage. Finally P refers to the results of the t-test, with P the probability that the means of the distributions of actual and modeled delay are equal. Any probability greater than 0.05 indicates there is no evidence that the means are different.

Table 13. MIT spacing model calibration

Airport	$\mu_{\varepsilon\text{-}MIT}$ [%]	$\sigma_{\varepsilon\text{-}MIT}$ [nm]	D_{act} [min]	$\mu_{\varepsilon\text{-}D}$ [min]	$\sigma_{\varepsilon\text{-}D}$ [min]	$\mu_{\varepsilon\text{-}D}$/ D_{act} [%]	P
PHL	5%	2	1.5190	0.0381	4.0295	2.51%	0.7237
LGA	5%	2	1.7002	-0.0056	4.6085	-0.33%	0.9695
EWR	-10%	2	2.0969	0.118	4.9453	5.33%	0.5835
JFK	25%	2	3.8469	-0.1656	7.8363	-4.30%	0.7915
TEB	15%	2	1.9406	-0.0206	6.0665	-1.03%	0.9691

The results in Table 13 suggest that at most airports the separations applied by controllers is greater than that specified. Only at EWR are the applied separations lower than specified. According to interviews, however, applied separations are generally lower than specified, as suggested by the results for EWR. This suggests that some other factors may be influencing the higher calibration separations. No further analysis into these other factors has been completed, due to time constraints.

Each airport was calibrated independently, to allow for modeling of different procedures at different airports. Ideally, each specified MIT (10 MIT, 15 MIT, 20 MIT etc.) would also be calibrated independently for each airport. However, due to time constraints this was not completed, and may be completed in future work.

Flights unaffected by MIT Restrictions Metering

When no restrictions are applied to a flight, safety separation requirements are 5 MIT in ARTCC airspace. However, according to expert elicitation from TMCs at ZNY, and according to McTMA researchers, most controllers separate aircraft by 7 MIT instead, to ensure that this separation is not violated. A separation standard of 7 MIT is thus applied in the modeling of the baseline instead of 5 MIT separation.

Even if a flight is not affected by MIT it may still not arrive at its ETA. This is because the ETA represents an estimate of the arrival time given the aircraft's entry time into the system, and does not represent its actual arrival time. An error representing deviation from the ETA is thus added for aircraft that are not delayed by MIT upstream of the TRACON, and are thus modeled to arrive at their ETA at the TRACON boundary. This error is generated for each tier through which the aircraft flew unimpeded. The error is sampled from a normal distribution centered at zero and with a standard deviation of 90 seconds. A 90 second standard deviation was chosen based on the average standard deviation of the distributions of unimpeded transition time, which was calculated to be 88 seconds.

4.2. TRACON Delay

Even when the MIT restrictions upstream of the TRACON are sufficiently high that the TRACON arrival rate is lower than the airport capacity, aircraft may still be delayed in the TRACON to ensure that runway separation requirements are met. This is because the arrival flows entering the TRACON from different arrival fixes are not coordinated, and aircraft from different arrival fixes need to be slotted together in the final approach to the landing runway/s.

TRACON Delay Model

Figure 47 shows the TRACON delay model that takes as input the ATAs at the TRACON boundary that are the output of the delay flow model shown in Figure 27.

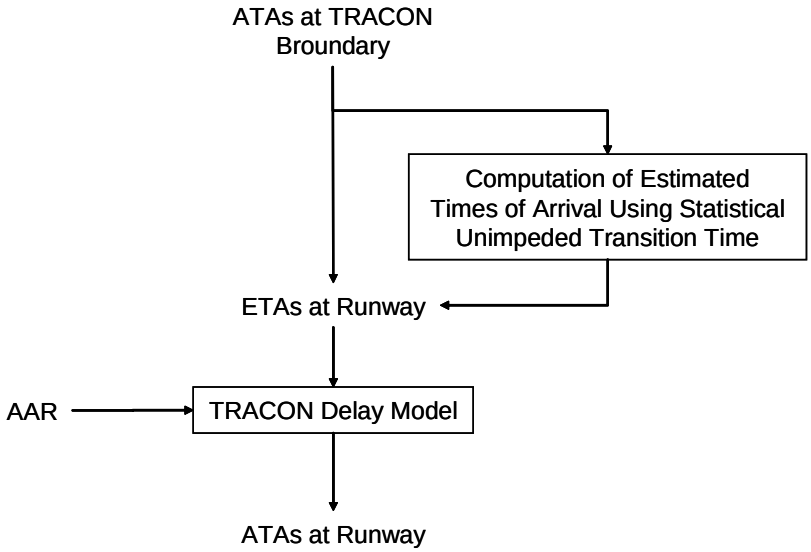

Figure 47. TRACON delay model diagram

Unimpeded transition times between the TRACON boundary and the runway are added to the ATAs at the TRACON boundary to derive ETAs at the runways. Then a TRACON delay model adds any needed delay to each flight such that the AAR at the runway is satisfied. The TRACON delay is added to the runway ETAs to compute runway ATAs for each flight.

The delay incurred in the TRACON is modeled as a function of the runway capacity limitations. As described in Section 3.1.1 the airport AAR describes the overall runway configuration capacity, although it is generally under-described by the reported AAR. Instead an AAR between the reported AAR and the 100^{th} percentile of the throughput observed under that reported AAR is generally the best description. The AARs chosen to constrain the arrival flows in the TRACON were thus used to calibrate the TRACON model against actual TRACON delays. The TRACON model calibration is described in detail below.

The arrival flows into the airport are metered to meet the applied AAR using time slots, in a similar way to the use of time based metering by McTMA. The TRACON

model is different to the McTMA time base metering model, however, in that the allocation of time slots is applied only in the TRACON, and only after the aircraft are metered by MIT upstream of the TRACON. All the delay is absorbed in the TRACON, and none is fed back upstream, as in McTMA time based metering. By calculating each aircraft's delay in order to fit it into a time slot calculated from the flight's ETA, the demand from other traffic, and the applied AAR, each flight's expected TRACON delay is estimated.

TRACON Delay Model Calibration

As described above the TRACON delay model was calibrated by adjusting the AARs chosen to constrain the arrival flows in the TRACON, so as to equate the distributions of actual and modeled delay in the TRACON as closely as possible. The AAR was adjusted by applying different percentiles of the capacity envelopes developed for each airport, for each reported AAR, some of which are presented in Figure 8 to Figure 12. The TRACON model was calibrated by reducing the mean difference between the actual and modeled delay within the TRACON. Further calibration using other parameters would be required to reduce the difference between the standard deviations of the actual and modeled delay distributions at the TRACON boundary. However, as in the calibration of the MIT spacing model, this is a secondary effect, and was not completed due to time constraints.

The results of the TRACON delay model calibration are presented in Table 14. The percentile of the capacity envelopes that calibrated best for each airport is included in the table. The t-test was completed on the results of the calibration to measure the performance of the calibrated model against actual delays. The t-test measures the probability that the means of the distributions of actual and modeled delay are equal.

In Table 14 $\text{Percentile}_{\text{CapEnv}}$ refers to the percentile of the capacity envelope that calibrated best. D_{act} refers to the average actual delay in the TRACON calculated according to actual transition times relative to estimated unimpeded transition times. $\mu_{\varepsilon\text{-}D}$ refers to the mean error between the modeled delay and the actual delay D_{act}. Similarly $\sigma_{\varepsilon\text{-}D}$ refers to the standard deviation of the error between the modeled delay and the actual delay D_{act}. The ratio of the mean error $\mu_{\varepsilon\text{-}D}$ and the actual delay D_{act} is also presented as a percentage. Finally P refers to the results of the t-test, with P the probability that the means of the distributions of actual and modeled delay are equal. P values greater than 0.05 indicate that there is no evidence that the means are different.

Table 14. TRACON delay model calibration

Airport	$\text{Percentile}_{\text{CapEnv}}$ [%]	D_{act} [min]	$\mu_{\varepsilon\text{-}D}$ [min]	$\sigma_{\varepsilon\text{-}D}$ [min]	$\mu_{\varepsilon\text{-}D}$/ D_{act} [%]	P
PHL	91%	3.5031	-0.0131	9.6961	-0.37%	0.9285
LGA	88%	3.5133	0.1949	8.5602	5.55%	0.1109
EWR	88%	1.9129	-0.1777	6.8739	-9.29%	0.0747
JFK	95%	1.2154	-0.1422	5.6645	-11.70%	0.1733
TEB	96%	0.4702	-0.0719	4.1549	-15.30%	0.4767

The results in Table 14 suggest that LGA, EWR and PHL are more constrained relative to the percentiles of the capacity envelopes observed than JFK and TEB. This is expected as JFK and TEB do not show the same degree of saturation as the other airports. The maximum throughput observed at these airports, from which the 99th percentile capacity envelopes were developed, is thus expected to be closer to the applied constraint on these airports under baseline operations, as identified in the TRACON delay model calibration.

The TRACON throughput under the TRACON model was also compared to that under actual operations. Table 15 below presents the throughput measured per quarter hour under actual operations (τ_{Actual}), and under modeled TRACON operations (τ_{TRACON}), and the difference between the two. It is clear that the differences are very low – less than 1 % of the actual TRACON throughput in all cases except TEB. At TEB the percentage is higher primarily because the throughput at the airport is lower.

Table 15. TRACON throughput comparison under baseline operations

Airport	**Percentile$_{CapEnv}$** [%]	**Average τ_{Actual}** [ac/hr]	**Average τ_{TRACON}** [ac/hr]	**Average (τ_{Actual}-τ_{TRACON})** [ac/hr]	**Average (τ_{Actual}-τ_{TRACON}) / τ_{Actual}** [%]
PHL	91%	25.3	25.1	0.17	0.68 %
LGA	88%	24.2	24.2	0.01	0.06 %
EWR	88%	25.4	25.3	0.14	0.56 %
JFK	95%	17.9	17.9	-0.02	-0.09 %
TEB	96%	9.9	10.2	-0.29	-2.91 %

As in the MIT spacing model calibration, each airport was calibrated independently, in order to model different procedures, and TRACON capacities for the different airports. This is despite the fact that EWR, LGA, JFK and TEB are all in N90. This is because, although each airport is in the same TRACON, the arrival flows are still separate, and the delay able to be absorbed is thus independent.

Error in Meeting TRACON arrival slots

An error is added to flight arrival times at the runway to model TRACON controller inefficiencies at meeting runway arrival slots. Even if a flight is not metered within the TRACON, it may still not arrive at its ETA, because the ETA represents an estimate of the arrival time given the aircraft's entry time into the system, and not its actual arrival time. An error representing deviation from a flight's ETA or arrival slot is thus added to all flights. The error is sampled from a normal distribution centered at zero and with a standard deviation of 90 seconds. A 90 second standard deviation was chosen based on the average standard deviation of the distributions of unimpeded transition time, which was calculated to be 88 seconds.

4.3. Holding Procedures

When aircraft must be delayed by an amount that is too great to be absorbed by vectoring or speed control, the aircraft must be put into a holding pattern. Holding is, however, generally avoided to the extent possible, because it is highly restrictive. Holding is restrictive because the accuracy with which aircraft can be released from the holding pattern to meet a specified aircraft separation is greatly reduced. This is because there is significant variability in where in the holding pattern the aircraft could be, and in which direction it is flying, when released. As mentioned above, N90 and ZNY do not hold aircraft for N90 or PHL, while the surrounding facilities are able to hold aircraft more effectively. There is however still limited holding capacity near the N90 boundary in ZDC, and the amount of traffic from ZBW, where aircraft can hold near the TRACON boundary, is considerably less than that from ZNY and ZDC. Consequently, a managed reservoir such as is used at PHL TRACON, where aircraft are held close to the TRACON boundary to provide a reservoir from which arrivals can be pulled whenever the runway is idle, is not used in N90. In N90 the restrictions passed back are instead kept low enough to ensure pressure on TRACON airspace and runways is maintained, while being high enough to ensure that the delay needed to be absorbed can be absorbed.

When a facility is unable to cope with any more aircraft, because the flow into the facility has not been metered sufficiently enough by the restrictions in place, no-notice holding may be specified to the upstream facility. This requires the upstream facility to hold all aircraft inbound into that facility, without any pre-planning, and can be difficult and disrupting. No-notice holding is a standard procedure in PHL TRACON, where arrivals are refused and forced to hold in the surrounding facilities. However, no-notice holding is less common in N90, because of the limited opportunity to hold aircraft near the TRACON boundary.

Baseline Modeling of Holding

In the baseline model holding is assumed to be a component of MIT and TRACON metering in that holding delays aircraft so that they will meet either the specified MIT restrictions, that they will not violate safety separations requirements en-route or in the TRACON, and that they will not violate runway acceptance rates. By constraining the baseline model with MIT, safety separations requirements, and runway acceptance rates, as described in the sections above, holding does not thus have to be modeled explicitly. The absorption of the required amount of delay is modeled, irrespective of the method by which delay it is absorbed, be it holding, vectoring, or speed control.

Modeling holding procedures explicitly would require modeling the causes of going into holding, the number of aircraft able to be held, and release from the holding pattern. The causes for holding may include ones other than runway and TRACON capacity constraint such as local enroute weather and congestion. Such holding procedures were considered irrelevant to McTMA and were therefore not modeled.

4.4. Internal Departure Release Procedures

Merging internal departures into N90 and PHL arrival flows is particularly important in ZOB, ZBW and ZDC, as each of these facilities have significant numbers of internal departure to N90 and PHL[6]. After identification of a gap in the overhead flow, TMCs specify release times to internal departures which allow them to hit the gaps in the streams at specified merge points[7]. This requires estimates of the transition times from the airport to the merge point, and the speed with which the gap in the stream is moving. The exact procedure by which release times are specified differs from facility to facility. TMCs in ZBW make use of tables of times from airports to fixes to specify the required release time, while TMCs in ZOB and ZDC make use of predetermined release points from which it takes the same amount of time for an en-route aircraft to fly to the merge point, as it does an internal departure to fly to the merge point from its departure airport. This is illustrated in Figure 48. Due to the inaccuracy of the estimates, and to avoid loss of capacity, as any gaps created cannot be closed, TMCs aim to release aircraft just ahead of gap in the flow. This forces the aircraft to be slowed down in the air in order to slot into the gap in the stream for which it was released, incurring unnecessary delays.

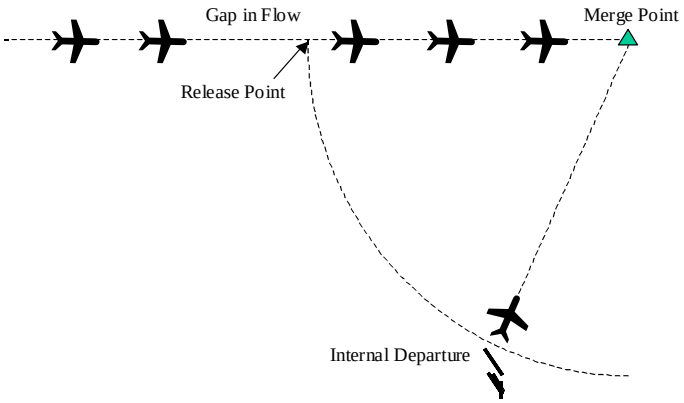

Figure 48. Internal departure release point

Baseline Modeling of Internal Departure Release

Internal departure release was modeled in the baseline model by allowing the absorption of delay at those internal airports identified as significant (more than approximately 1% of operations to the destination airport), so as to meet a specified MIT restriction. Any delay required to be absorbed to meet a restriction was thus assumed to be absorbed on the ground. MIT and safety separation requirements at any merge points

[6] Internal departures are released using DSP (Departure Spacing Tool, by Computer Sciences Corporation), which coordinates departures from a number of airports. Departures are put on APREQ (Approval Request) meaning that each departure must request the TMC at the Center to specify a release time. This time includes a three minute departure window in which the internal departure must depart.

[7] Each TMC may develop a different release point on the radar screen at which he/she releases aircraft for each airport.

downstream of the internal airport were also enforced, thus modeling delay incurred so as to fit into the overhead stream.

4.5. Miscellaneous Observations

Another limitation that was observed in the site visits is the possibility of gridlock effect between interacting flows in an airspace. An example was given in Section 3.1.3 where ZOB holds the N90 departures from ZNY when ZNY holds the N90 arrivals from ZOB. ZOB holds the N90 departures because the proximity between the arrival and departures airways makes the arrival holding pattern block the departure airways. Because the arrival and departures are interdependent, where departures need to leave the airports in order to make room for the arrivals to land, a gridlock effect is created where both arrivals and departures are holding and waiting each for the other to advance. Gridlock rarely materializes because the system reacts in time to prevent its extremely adverse consequences. As the system approaches gridlock, however, it is expected that the throughput of the affected resource be reduced. Because gridlock effects are approached as congestion is increased, the effect is expected to be similar to that of the controller workload increase depicted in Figure 5. In other words a reduction in the resource throughput at extremely high congestion and delay levels.

Another observation that was made during the site visits is the lack of en-route delay measurement. Delays are in general reported by controllers (using flight strips or other mechanisms) and documented by traffic managers in traffic management logs only if above 15 minutes in any sector. This usually occurs only when an aircraft is held as all delays without holding are significantly below this limit (see Section 3.1.3). However, as aircraft incur below 15-minute delays in successive sectors the total accumulated delay may be significantly above 15 minutes, none of it is reported. As a result the traffic managers have no notion of arrival delays except in the case of holding. As a feedback about the delay level in the system they use the departure delays at the major airports. These are also reported only if above 15 minutes. The traffic managers attempt to balance arrival and departure delays in order to balance the arrival and departure flows and avoid the gridlock effects that may result from imbalanced delays. Therefore, they monitor the departure delays reported, and once above a certain limit they increase arrival delays (through holding) to match the departure delays. However, with the lack of accurate arrival delay measurement, the effectiveness of this procedure is questionable. One possible use of McTMA, which is not yet addressed in this work, is to provide a more accurate measurement of arrival delays at each meter fix along the route to destination.

Another observation that was noted during the site visits is the emphasis of the Traffic Managers on the operations during thunderstorms. The most disrupting event in the NAS is usually thunderstorms that are localized and affect certain arrival gates, at times of high demand (typically in the summer). One indication was made by a TMC to the possible use of McTMA to suggest rerouting of and assigning aircraft to the appropriate fixes (open gates) during thunderstorms. This function is not currently explicitly defined for McTMA (in the next section) but may be possible to capture through balancing of arrival flows and will be investigated with NASA McTMA researchers.

5. Identification and Modeling of McTMA Functionality

In order to identify and quantify the benefit mechanisms of McTMA, the core functions of McTMA were identified and modeled. Based on review of TMA and McTMA literature [2, 3, 4] and discussions with NASA domain experts, the following functions of McTMA were identified:

1. Time based metering, with “delay feedback” and “capacity distribution”
2. Dynamic metering
3. Tiered metering
4. Demand visualization
5. Multiple facility coordination
6. Internal departure scheduling
7. Runway assignment (according to NASA’s McTMA researchers this function is not currently implemented although it is described in [4])

Some of these functions are functions of both TMA and McTMA (such as time based metering, dynamic metering, demand visualization, internal departure scheduling and runway assignment), while others are specific to McTMA (such as capacity distribution, multiple facility coordination and tiered metering). Each function is described in detail below, in the context of the queuing abstraction of the arrival flow into N90 and PHL described in Section 2.2. This abstraction provides a means by which the effect of each function can be analyzed, and benefit mechanisms identified accordingly. Key benefit mechanisms and resulting quantifiable benefits are presented for each function in Section 6.

While benefit mechanisms were derived for all the functions, only the most important functions (based on NASA’s feedback) were modeled and their benefits assessed quantitatively, due to time constraints. The functions that were ultimately modeled are:

1. Time based metering, with “delay feedback” and “capacity distribution”
2. Dynamic metering
3. Tiered metering
4. Multiple facility coordination
5. Internal departure scheduling

5.1. Time Based Metering with Delay Feedback and Capacity Distribution

Function description

Time based metering is identified to be the core functionality of McTMA that generates scheduled times of arrival (STAs) at meter fixes for which the McTMA system is implemented. These STAs are calculated from estimated times of arrival (ETAs) at the meter fixes. The calculation of ETAs assumes that the aircraft are unimpeded (as

described in Section 3.2). STAs differ from the ETAs according to the amount of metering required to ensure efficient use of the airspace, and is based on satisfying sequence and schedule constraints, satisfying the required arrival rate at the downstream meter fix (or runway), and the feedback of delay from downstream. McTMA also attempts to maintain a first come first serve (FCFS) sequence, based on the ETAs. By doing this, STA deviation from the ETAs is kept to a minimum. STAs are never specified earlier than their corresponding ETAs, only at the same time or later.

The scheduling constraints associated with time based metering include [4]:

- Meter fix acceptance rate
- Gate acceptance rate
- TRACON acceptance rate
- Super stream class separation distance (MIT)
- Meter fix blocked interval
- Airport acceptance rate
- Runway acceptance rate
- Runway occupancy time
- Wake vortex separation
- Runway blocked interval

The resource with the most constraining of these constraints is the bottleneck in the system, and defines the system service rate, and thus throughput.

Two key components in the calculation of STAs from ETAs and arrival rate constraints are “delay feedback” and “capacity distribution”. Delay feedback is the distribution of delay between meter fixes to maintain pressure on the bottleneck (often the runways) without exceeding the delay absorption capacity (delayability) associated with each resource. Delay feedback is accomplished by the feeding back of delays upstream, through meter fixes, when the downstream resource reaches its capacity limit to absorb delays. Capacity distribution is the process of generating arrival rate profiles (acceptance rate during successive time periods) for upstream meter fixes, according to STAs propagated back from downstream. The STAs at each meter fix, which correspond to downstream STAs, are counted to generate acceptance rate profiles for meter fixes immediately upstream. These acceptance rate profiles are then used to generate STAs for the upstream traffic. Using capacity distribution an aircraft is not forced to meet its STA that is propagated from downstream if not able to or inefficient, rather another aircraft may be substituted in its slot to avoid missing the slot, as long as the acceptance rate profile is not violated.

Time based metering enables delay feedback and capacity distribution through McTMA’s specification of the amount of delay to be absorbed in the airspace between meter fixes and adjusting the scheduled meter fix arrival times (STAs) accordingly.

Also included in time based metering is the specification of a freeze horizon, which is a set time period from the current time, for each meter fix, within which the schedule is fixed and cannot be changed. It is within this freeze horizon that all the delay required to meet the STA must be absorbed. The freeze horizon at DFW is currently 19 minutes, and those at PHL and N90 are expected to be similar. The freeze horizon is

designed to be long enough to ensure that the controllers working the aircraft scheduled at the meter fixes have enough time to absorb the delay required to meet the STA, but short enough to ensure that the STAs remain dynamic enough to accommodate changing constraints. The freeze horizons for different meter fixes in McTMA are not likely to be the same size. Particularly in some cases they may be less than 19 minutes.

Function modeling

Modeling time based metering requires definition of the flow network, and the meter fixes where the STAs will be generated. The definition of the flow network is highly specific to the airspace being modeled. The flows were identified for PHL and N90 through expert elicitation from NASA and at the facilities included, as detailed in Appendix A and Section 3.2 (Figure 49 shows as an example the flow network for JFK). Also required is the definition of the system boundary relative to the points on each flow at which metering will start. The system boundary is dependent on the extent of McTMA's airspace coverage, which is also related to the length of the time line of McTMA.

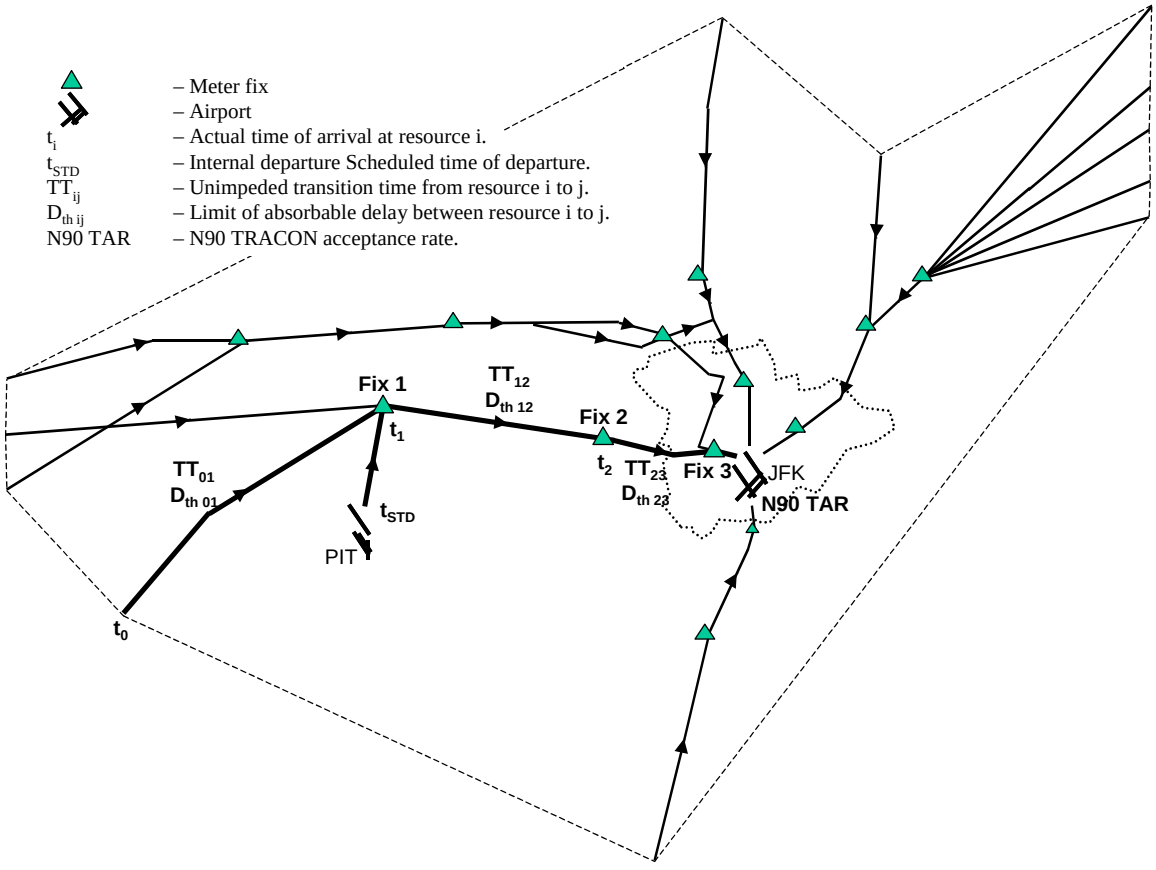

Figure 49. Network of meter fixes with system boundary and model inputs indicated

Other inputs to the modeling of time based metering include flight arrival times at the outer boundary of the system, departure times for internal departure (departures from

within the system boundary), unimpeded transition times between meter fixes, constraints, and the acceptable level of absorbable delays between each pair of meter fixes. This is illustrated for a single flow into JFK in Figure 49.

ETAs, STAs and actual times of arrival (t) are calculated according to a number of steps, as illustrated in Figure 50, and detailed step-by-step following. In this way ETAs, STAs and actual times of arrival are calculated at all applicable points, for all flights.

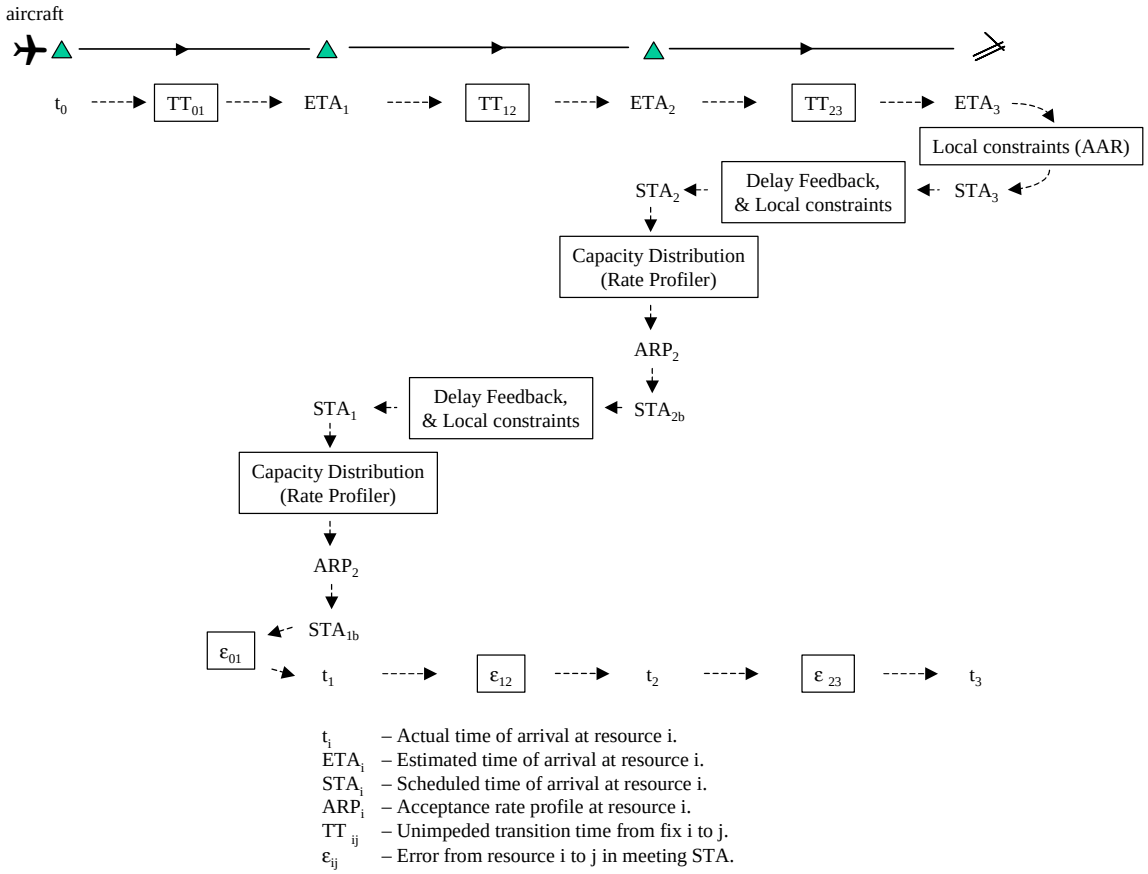

Figure 50. Calculation of STAs and actual times of arrival t_i

1. All flight ETAs are calculated at all applicable points (this includes all metering points on the aircraft's route downstream of the aircraft's current position). As detailed in Section 3.2, this requires as inputs each flight's initial conditions, unimpeded transition times to the downstream metering fix in the flight plan (TTs), and unimpeded transition times between subsequent points. As illustrated in Figure 50, in this manner each flight's ETAs are calculated for all the applicable points, working downstream from the system boundary to the runway threshold.
2. Each flight's STA is calculated at the runway according to the TMA algorithm described in the Dynamic Planner document by Gregory Wong [4]. This algorithm, referred to as f_{TMA} in Equation (6), accounts for local scheduling constraints listed

above, including runway acceptance rates, TRACON acceptance rates and separation requirements for safety.

$$STA_{rwy} = f_{TMA}\left(ETA_{rwy}, \text{local scheduling constraints}\right) \quad (6)$$

The local scheduling constraints in Equation (6) when applied at the runway included only the runway acceptance rate. The runway acceptance rate capacity is estimated as the limit throughput observed using historical data and is modeled as a function of the departure rate as presented in the capacity envelopes in Section 3.1. It was assumed that the runway acceptance rate capacity is enough to ensure the safe separation between aircraft. Specific separations between pairs of aircraft would require runway assignment knowledge, which was not modeled and not assumed as a function of McTMA as described above.

3. The STAs at the upstream meter fixes *i* are then calculated according to the STAs at the runway or downstream meter fix *i*+1; the unimpeded transition time $TT_{i,i+1}$ between the meter fix under question and the downstream runway or meter fix; and the feedback of delays from the downstream runway or meter fix, as given in Equation (7). The feeding back of delays from the downstream runway or meter fix, given by function f^n in Equation (7), is based on the threshold of delay that can be absorbed between the meter fix and the downstream runway or meter fix ($D_{thi,i+1}$), and the amount of delay that was not able to be absorbed in the downstream airspace (Delay Feedback in Equation (7)). This function is such that if more delay must be absorbed than the delay threshold, only the delay threshold is allocated for absorption at the fix under question (where the delay threshold is computed as described in Section 3.1.3). The rest of the delay is fed upstream to the upstream meter fix.

$$STA_i = STA_{i+1} - TT_{i,i+1} - f^n\left(D_{th\,i,i+1}, \text{Delay Feedback}\right) \quad (7)$$

4. Through capacity distribution and the rate profiler process, an acceptance rate profile is generated at each meter fix *i,* according to the STAs calculated in step 3. This is done by counting the number of STAs in bins using the sliding window described in the Dynamic Planner document by Gregory Wong [4]. These acceptance rates in addition to the safety separation requirements become the local scheduling constraints in Equation (6) applied at meter fix *i* in step 5.

5. New STAs are then calculated at each meter fix *i* using the TMA algorithm, Equation (6) applied at meter fix *i:* $STA_i = f_{TMA}\left(ETA_i, \text{local scheduling constraints}\right)$, according to the ETAs at meter fix *i* and the local scheduling constraints at meter fix *i*. These local scheduling constraints include the arrival rate profile and the safety separation requirements (7 Nautical Miles), both of which the STAs are ensured not to violate. These STAs may specify a different sequence to the earlier STAs calculated in step 3 based on the delay feedback.

   In this manner each flight's STAs are calculated for all the metering points, working upstream from the runway threshold to the delay boundary.

   As mentioned in Section 3.2 unimpeded transition times between meter fixes TT_{ij} are estimated through analysis of historical track data. Similarly, the delay thresholds D_{th}

between tiers are estimated based on expert elicitation and data analysis as mentioned in Section 3.1.3.

Both the delay threshold and the capacity constraints are model parameters that are varied to represent the effect of different benefit mechanisms that will be described in Section 6.

6. Actual times of arrival at the meter fixes and at the runway threshold (t_i and t_{rwy}), are calculated according to the STAs at the meter fixes, and an error due to variability in meeting the STAs specified. In time based metering controller imperfection in advising the aircraft to meet their STA lead to the actual time of arrival at meter fixes differing from the specified STAs by an error, ε. Thus, the actual time of arrival is calculated as follows in Equation (8):

$$t_i = STA_i + \varepsilon_i \quad (8)$$

The error ε was identified from previous NASA TMA field research to be centered around zero, with a standard deviation of 150 seconds. However, based on consultation with NASA's McTMA researchers 150 seconds was considered excessive and the error ε was modeled as a normal distribution centered around zero, with a standard deviation of 90 seconds, and maximum and minimum of two standard deviations.

From the actual times of arrival at the meter fixes and runway threshold, the system throughput (τ) and aircraft delay are calculated. These are key parameters in the benefit analysis. The system throughput is calculated as the average resulting flow rate at the final resource, which is calculated according to the actual times of arrival at the runway threshold.

Delay at each meter fix, and at the runway threshold (D_i and D_{rwy}), is calculated as the difference between the actual time of arrival at the meter fix and the corresponding ETA as shown in Equation (9).

$$D_i = t_i - ETA_i \quad (9)$$

5.2. Dynamic Metering

Function description

Dynamic metering refers to the dynamic nature of McTMA in which STAs are recalculated every 12 seconds, up to the freeze horizon (typically 19 minutes), after which the STAs become fixed. This is in contrast to current operations where a MIT restriction generally remains unchanged for the duration of the restriction, which may extend to few hours.

Function modeling

While STAs were not updated every 12 seconds, the dynamic nature of McTMA's time based metering was modeled by ensuring that STAs are calculated to meet variable acceptance rates and to match variable demand rates. It was also captured

by updating the STAs between upstream tiers and downstream tiers adjusting for accumulated errors in meeting the upstream STAs. The error in meeting the STA was added to each STA, at each tier, limited so as not to allow the ATA at the fix to be earlier than the ETA. Also, because McTMA updates ETAs and recalculates STAs every tier, it is able to adjust for the errors incurred in upstream tiers. This was modeled by sampling an error from the error distribution, adding it to the STA of the tier in which the error was incurred, and then subtracting it from the amount of delay to be absorbed in the tiers downstream. This was done because McTMA would adjust the STAs for the tiers downstream to account for the error upstream. The change in the amount of delay to be absorbed in each downstream tier was limited in that the resulting adjusted amount of delay to be absorbed must be greater than zero, and less than the delayability of the tier.

5.3. Tiered Metering

Function description

Tiered metering refers to metering at a number of meter fixes on each arrival flow, extending a number of tiers upstream from the destination airport. STAs are to be specified at each of these meter fixes. Current TMA implementations only meter at the TRACON arrival fixes, but McTMA is to meter at as many as 3 tiers of meter fixes. Tiers are carefully located according to where they would be most useful, such as at facility boundaries, merge points and arrival fixes. The tiers are expected to be approximately 100 nm, 200 nm, 300 nm and 400 nm from the TRACON boundary, although there will be significant deviation from these arcs. The exact location of the meter fixes is still under development, and is discussed in greater detail in Appendix A. Regional Metering is an extension of this concept, which extends time based metering further upstream to the entire en-route environment.

Function modeling

Tiered metering is modeled in the flow network to which time based metering is applied, as discussed above. The location of each meter fix is significant to the estimate of a number of the parameters in the model including the delay threshold D_{th}, transition times TT between meter fixes, and freeze horizons. STAs are propagated through merges and splits according to delay feedback and capacity distribution.

5.4. Demand Visualization

Function description

Demand visualization refers to the function of McTMA whereby demand can be visualized on loadgraphs and timelines. Loadgraphs show demand according to ETAs and STAs (pre-metered demand and metered demand), capacity, and expected delay. The timelines show each aircraft's ETA, and STA. Using the demand visualization function TMCs may use McTMA as a what-if analysis tool in order to assist in selecting more optimal MIT restrictions.

Function modeling

Demand visualization will not be modeled explicitly. Instead each mechanism by which benefits are realized from demand visualization will be modeled explicitly. It was planned to include modeling a what-if capability for McTMA through the use of an algorithm to select optimal MIT restrictions as discussed in detail in Section 6.7. However, after consultation with NASA's McTMA researchers (at the second phase briefing), it was decided not to model the what-if functionality of McTMA since it is no longer planned for, neither as a permanent function, nor as a transitional phase function in the McTMA deployment. Some benefit mechanisms enabled by the demand visualization are modeled as described in Section 6.

5.5. Multiple Facility Coordination

Function description

Multiple facility coordination refers to the coordination between all facilities involved in the arrival process that is enabled by McTMA. This includes the TRACON, surrounding Centers, and Centers adjacent to these Centers, if involved in the arrival process. For New York TRACON (N90) and Philadelphia TRACON (PHL), the associated centers are New York ARTCC (ZNY), Boston ARTCC (ZBW), Washington ARTCC (ZDC) and Cleveland ARTCC (ZOB). Atlanta ARTCC (ZTL) and Indianapolis ARTCC (ZID) are also involved in the arrival process but are not considered at this stage because of the cost limitations on the initial implementation of McTMA. With the use of McTMA all these facilities coordinate their arrival flows so that aircraft do not arrive at merge points with flows from other facilities at exactly the same time, but instead slot together with a minimum of resequencing. This leads to benefits, a number of which are discussed under other functions. The benefit mechanisms are detailed in Section 6.8.

Function modeling

Multiple facility coordination is modeled implicitly to time based metering and tiered metering as STAs are generated for flows merging from different facilities and transitioning between upstream and downstream facilities. Modeling of specific benefit mechanisms of facility coordination is discussed in Section 6.

5.6. Internal Departure Scheduling

Function description

McTMA specifies release times for departures that are internal to each Center in which McTMA is installed, and are destined for the TRACON for which McTMA is implemented. This ensures that these internal departures are included in the time based metering of the arrival traffic into the TRACON. However, because they are not en-route they are scheduled into the arrival flow before they take off. They are allocated STAs, like any other flight, but absorb the delay for the first meter fix on the ground, before take off. This means that delay need not be absorbed in the limited airspace between the departure airport and the first meter fix.

One major complication dealing with internal departures is the large uncertainty in their scheduled departure time and estimated time on the ground. This large uncertainty reduces the accuracy of estimating the demand on the TRACON and runway system for metering and for deciding on landing slots for the internal departures. This issue remains a challenge and a research area for McTMA.

Function modeling

In the McTMA model, internal departures are included in the generation of ETAs, and assignment of STAs and departure release times. As described in Section 3.2, internal departures enter the system at their actual takeoff time (measured by the first track) rather than at their scheduled departure time. The ETAs are calculated based on the average time to arrive at the first meter fix after departure, and then subsequent meter fixes. The release time is then based on the assigned STAs at the downstream runway and meter fixes, and any delay required in addition to the actual takeoff time is absorbed on the ground. Therefore, for simplicity and conservatively, modeling internal departures here does not capture the impact of McTMA on the original takeoff time. This would require using scheduled departure times and estimates of the time spent on the ground, and dealing with their inaccuracies, which are still a subject of research for McTMA.

5.7. Runway Assignment

Function description

According to [4] McTMA includes a function that assigns runways to arriving flights by means of a runway assignment algorithm designed to reduce delay in the entire system. However, according to NASA's McTMA researchers this function is not currently implemented and the ATC operators may instead set runway assignment manually. Therefore, this function has not been modeled in this benefit assessment.

6. Identification and Modeling of McTMA Benefit Mechanisms

The benefits of applying McTMA to the PHL and N90 arrival flows are identified by applying the functions of McTMA identified in Section 5 to the constraints and flow management techniques of the PHL and N90 arrival flows identified in Section 3. First the benefit mechanisms identified in TO10 [6] and TO33 [7] are reviewed. Then a formal method for mapping McTMA functions into benefits is outlined and used in the derivation of the benefit mechanisms.

6.1. Review of Previous Studies

The benefits of TMA and McTMA were identified through the modeling of a few key benefit mechanisms in TO10 [6] and TO33 [7] respectively. These are described below:

TO10 [6] identified the following benefit mechanisms for TMA:

1. Time-Based Arrival Metering – Time based as opposed to distance based arrival fix metering.
2. Arrival Fix Delivery Accuracy – Improving the accuracy in meeting STAs.
3. Internal Departure Release – Internal departure delays are reduced by more efficient internal departure release into arrival gaps.
4. Center/TRACON Delay Distribution – An optimum balance between fuel burn savings, obtained through the absorption of delay in higher altitude ARTCC airspace; and runway system throughput, obtained through increased pressure on the runway by absorbing delay in the TRACON.

TO33 [7] identified the following benefit mechanisms for McTMA at PHL:

1. Demand Visualization – Visualization in time lines and load graphs allowing TMCs to plan to fully exploit an available arrival rate.
2. Coordination of Restrictions – The use of what-if functionality to test which MIT restrictions meet the available arrival rate, and balance delay and workload.
3. Shut-Off Decisions – Improved efficiency of shut-off and flow-resumption decisions
4. Off-loading to Reliever Runway – Identification of flights that can be off-loaded to a reliever runway

TO33 [7] studied three different concepts – a TRACON centric concept, which uses McTMA for demand visualization only, and does not implement time based metering; a fully dependent concept, in which time based metering is implemented completely; and a transition concept, which applies time based metering to one ARTCC flow, leaving the other ARTCCs to operated distance based metering. According to McTMA researchers McTMA is now to operate with full time-based metering, using the fully dependent concept only. In TO33 each of these concepts were modeled separately, and included the modeling of the appropriate benefit mechanisms described above. Not

all of the benefit mechanisms were modeled however. In particular, the off-loading to a reliever runway was not modeled.

TO33 [7] also identified the following operational limitations:

1. Availability of airspace for metering/vectoring – Small sector size leads to early holding.
2. Unavoidable delay – Not all delay can be addressed by CTAS.
3. Proximity of ZOB/ZNY boundary – Some delays must be incurred through MIT restrictions in ZOB.
4. Inherent inaccuracy of Miles-In-Trail approach – Sub-optimal selection of restrictions by TMCs.
5. Accuracy of compliance to restrictions – Error associated with meeting restriction.
6. TMC workload – Sub-optimal selection of restrictions by TMCs

No clear formal procedure was followed in TO10 [6] or TO33 [7] for the identification of benefit mechanisms. Instead the benefit mechanisms were described only qualitatively, making the mapping of functions to benefits unclear. This makes it unclear whether or not all benefit mechanisms were identified, implying that benefits may have been underestimated.

It was thus decided that a more formal, reviewable process is required for the identifications of benefit mechanisms, as described below in Section 6.2.

6.2. Analysis of Benefit Mechanisms

For the purpose of clarity, consistency, and completeness, functions, constraints, benefits, and benefit mechanisms were formally defined in Section 2.1. Based on these definitions, the benefits of McTMA are identified by applying the functions of McTMA identified in Section 5 to alleviate the current operations constraints and limitations identified in Section 2.3. A function excites a benefit mechanism by alleviating system constraints, which creates a benefit. This process results in a mapping from functions to benefits through benefit mechanisms. The mapping of functions to benefits is not one to one, nor are there a consistent number of steps in each benefit mechanism.

Charts of the mapping from functions to benefits through benefit mechanisms were developed to improve clarity and reviewability, and are included with the descriptions of the benefit mechanisms. A single chart is presented for each McTMA function, followed by a description of each mechanism enabled by the function under consideration. The following definitions apply to the charts:

- Bold shaded blocks represent McTMA functions.
- Normal shaded blocks represents characteristics of the function from which they are extracted.
- Bold unshaded blocks represent quantifiable benefits.
- Normal unshaded blocks represent intermediate steps between a function and a benefit, which make up McTMA benefit mechanisms.
- Black blocks with white writing represent direct economic benefit (used in Section 8).

- An arrow (to a function) represents “enables” (i.e. one function enables another function)
- An arrow (to a benefit or intermediate step within a benefit mechanism) represents “results in” (i.e. one function results in a benefit)
- Dotted arrows represent an effect that is not a current McTMA benefit, or benefit mechanism, but through enhancement of the functionality of McTMA, may become one. These benefits or benefit mechanisms shall not be modeled.
- Merging arrows mean that more than one function or benefit mechanism enables or results in a particular function, benefit or benefit mechanism.
- Forking arrows mean that more than one function, benefit or benefit mechanism is enabled by or results from a particular function, benefit or benefit mechanism.

When a modeling parameter can be related directly to a box in the charts, the parameter is included in parenthesis below the block.

The functions of McTMA for which benefit mechanisms are described are as follows:

1. Time based metering
2. Delay feedback and capacity distribution
3. Dynamic metering
4. Tiered metering
5. Demand visualization
6. Multiple facility coordination
7. Runway assignment
8. Internal departure scheduling

Although delay feedback and capacity distribution were described as part of the time based metering function in Section 5.1, because they are an integral part of the STA generation algorithm, they are considered as separate functions here, for the purpose of mapping to benefit mechanisms. Benefit mechanisms for each function are described in detail in Section 6.3.

6.3. Benefit Mechanisms of Time Based Metering

Time based metering is a function of McTMA and refers to metering of the arrival flow by specifying times of arrival at designated meter fixes. The mapping of the benefits of time based metering is presented in Figure 51.

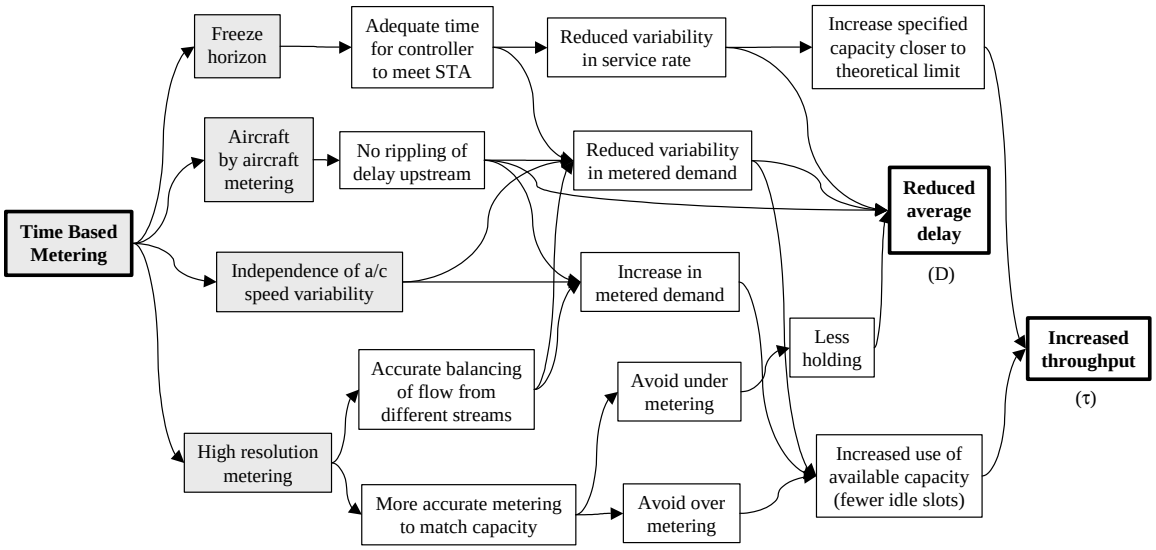

Figure 51. Benefits of time based metering

Time based metering has the following characteristics:

- Metering is aircraft by aircraft, and not pair wise
- Metering has high resolution
- Metering is independent of aircraft speed variability
- A freeze horizon is defined

The benefit mechanisms associated with each of these characteristics are presented on the following subsections.

6.3.1. Aircraft by Aircraft Metering

Aircraft by aircraft metering refers to the metering of individual aircraft independently, and not relative to the aircraft that they are trailing, as in current operations with MIT or Minutes in Trail (MinIT) metering. The benefit mechanisms of aircraft by aircraft metering are presented in Figure 52.

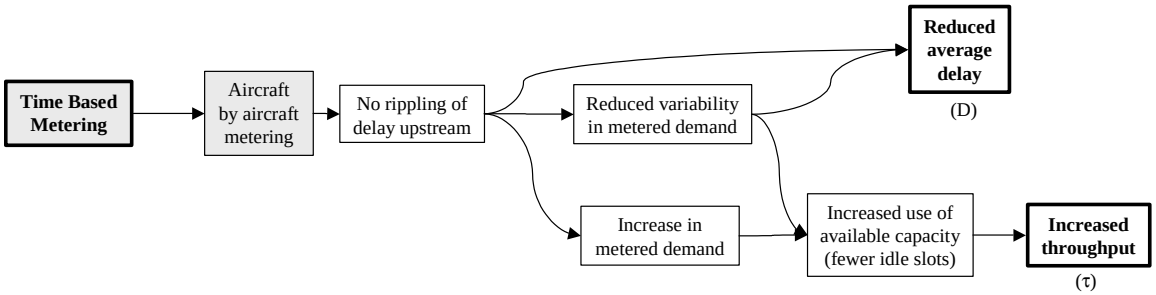

Figure 52. Benefit mechanism of aircraft by aircraft metering

Metering each aircraft independently means that any single aircraft's delay relative to its STA is not propagated, or rippled, upstream. Rippling occurs when a leading aircraft is delayed, and the trailing aircraft is metered relative to the delayed leading aircraft. It thus incurs the same delay as the first aircraft, as do all following

aircraft. In aircraft by aircraft metering delay is not rippled, as an aircraft's STA is not dependent on the leading aircraft, except to ensure the 5 MIT safety separation limit. Only if this safety separation were to be broken would delay be rippled upstream. (Alternatively, the delayed aircraft could be pulled from the schedule and rescheduled in the back of the sequence, avoiding any delay rippling.) Therefore, no rippling of delay upstream reduces average delay.

A reduction in the upstream rippling of delay results in increased metered demand because fewer gaps are opened in the arrival stream. Increased metered demand in turn results in increased use of the available capacity of the downstream resources as more pressure is applied on them. Applied to the constraining resource this results in fewer idle slots and thus increased throughput.

A reduction in the rippling of delay results in a more stable and robust schedule. This is expected to result in a reduction in the variability in the metered demand. Reduced variability in metered demand results in an increase in how much of the available capacity is used, as illustrated in Figure 53. The shading under the capacity in Figure 53 represents periods when demand drops below capacity and capacity is not fully used. These shaded regions are reduced when the variability in demand is reduced. Increased use of available capacity results in fewer idle slots in the flow, and thus, if applied to the constraining resource, increased system throughput.

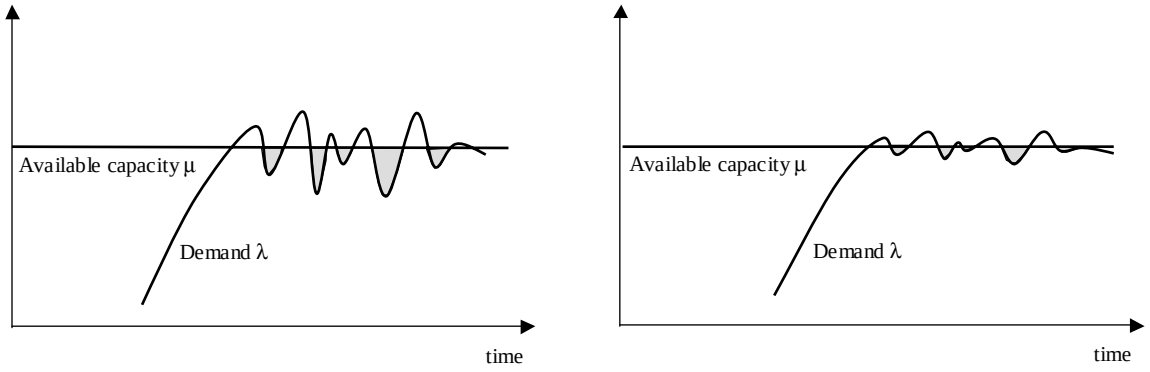

Figure 53. Effect of reduced variability in metered demand

Reduced variability in metered demand also results in a reduction in delay directly. This is illustrated in Figure 54. Variability can cause a high demand, followed by a low demand, averaging over the entire period to an average demand. If the service rate were equal to this average demand, the overall throughput would also be equal to the average value. However, those aircraft arriving when the demand was high would incur significant delays as illustrated in the upper part of the figure. If the aircraft arrived at the average demand rate in the first place, however, the throughput would still be average, but none of the aircraft arriving would incur significant delays, as illustrated in the lower part of the figure. The average delay is thus higher when the variability in demand is high, even if the overall throughput is maintained at the same level.

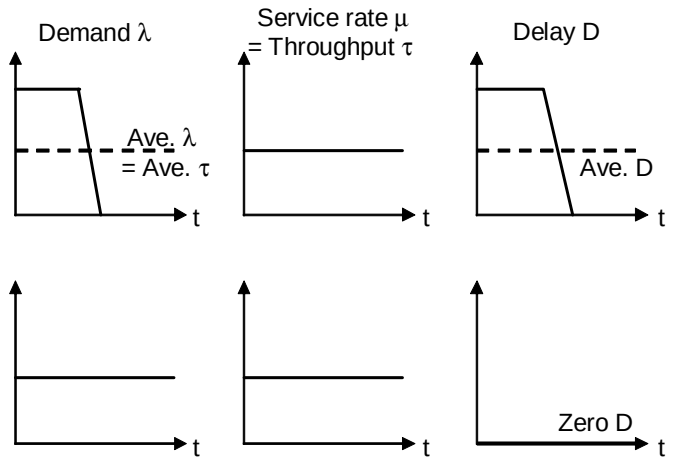

Figure 54. Higher delay resulting from high variability in service rate

Modeling of Aircraft by Aircraft Metering

Modeling of this mechanism is predominantly implicit to the modeling of time based metering described in Section 5.1, as aircraft STAs are specified without reference to the leading aircraft. In the modeling of distance based metering however, described in Section 3, STAs are calculated from MIT restrictions, relative to the leading aircraft's time of arrival at the meter fixes, and thus delays are rippled. Explicit modeling is however required for the controller action in time based metering, when a flight is late enough to infringe on the safety separation limit of the following aircraft. The trailing flight must then be delayed, or the late flight rescheduled. This was modeled by calculating the separation between aircraft at each meter fix, according to actual times of arrival (t). These were then adjusted if needed, delaying the trailing aircraft by the amount required to ensure that 7 MIT (a conservative separation used by McTMA instead of 5 MIT) safety limits are maintained. A 7-mile separation requirement is currently imposed at all fixes/tiers.

6.3.2. High Resolution Metering

Higher resolution metering than in current operations is a characteristic of time based metering, and refers to the resolution with which restrictions can be applied. MIT restrictions in current operations are specified in 5 MIT increments, which at a speed of 200 knots provide 1.5-minute time separations. McTMA calculates STAs to the nearest second. Although STAs are presented to the controllers to the nearest minute, more than one aircraft can be presented in the same minute slot, as long as the safety separation requirements are maintained. According to consultation with McTMA researchers, it is also reasonable to assume that the delay values shown in the metering lists given to controllers at final McTMA metering fixes (where little delay can be absorbed, for example, those in ZNY where 1-2 minutes can be absorbed at most) will be in tenths of minutes (for example 0.6 minute). The benefit mechanisms of high resolution metering are presented in Figure 55.

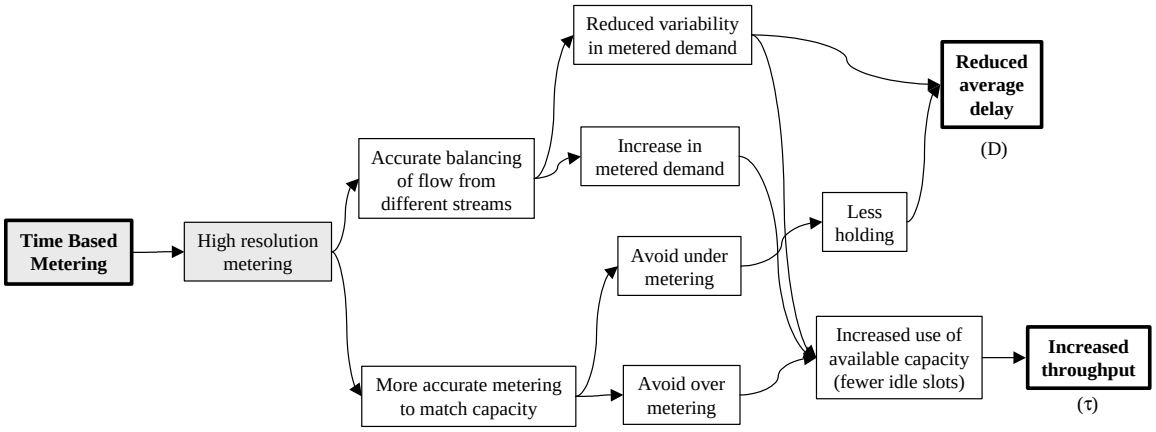

Figure 55. Benefit mechanism of high resolution metering

Higher resolution metering allows more accurate metering to match the available capacity. This is illustrated in Figure 56, which displays the metering rate (per hour) achieved in simulation applying different values of MIT and time based metering to a stream of traffic. The distance based MIT increments of 5 could only achieve discrete levels of metered rate. On the other hand time based metering was able to match any imposed metered rate (for example a rate of 27 arrival per hour in the figure), which was not achievable through MIT. (The plots in Figure 56 were generated using a simulation of distance based and time based metering applied to a Poisson arrival process). In the case of MIT metering (distance based metering), the resolution is low, and a lot of capacity can be lost. However, in the case of time based metering, the resolution is considerably higher, and consequently the flow can be metered to reduce the flow rate to almost exactly the correct level.

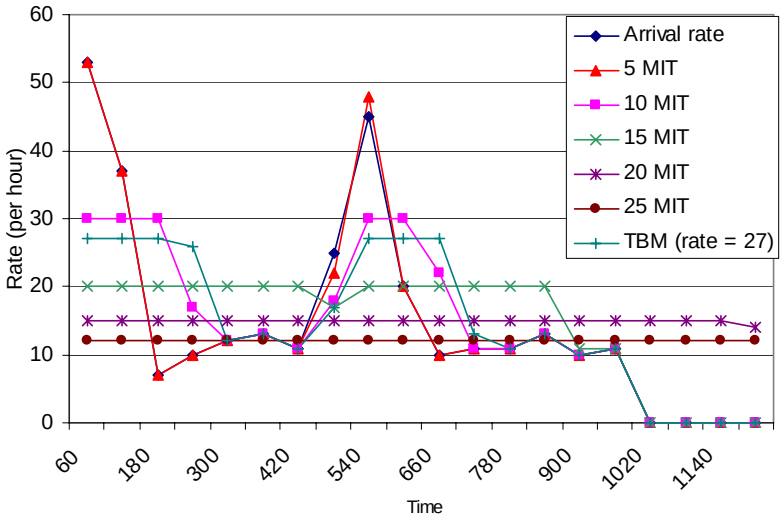

Figure 56. Resolution of distance based and time based metering

More accurate matching of capacity means that both over-metering (where metered demand is too low) and under-metering (where metered demand is too high) can be avoided. Avoidance of over-metering results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus

increased throughput. Avoidance of under-metering results in less holding (or other means to absorb the excessive queue) which results in reduced average delay.

Higher resolution metering also results in more accurate balancing of flows from multiple streams. Particularly merging flows can be metered more precisely according to the demand and to the constraints propagated upstream. Each individual aircraft is metered to fit into a gap in the stream into which it is to merge, reducing the gaps in the arrival stream and reducing the variability of the flow after the merge point. Reducing the gaps in the arrival stream results in increased metered demand, which in turn results in increased use of the available capacity of the downstream resources, as discussed in Section 6.3.1. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. As discussed in Section 6.3.1 reduced variability in metered demand also results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Figure 57 compares time based metering to distance based metering applied to two streams with imbalanced flows. In the first demand peak the demand was equally high from both streams. In this peak both 10 MIT and time based metering were able to meter the demand to the level imposed (the horizontal line). The second demand peak was caused by high demand from only one of the two streams. At this peak, time based metering was able to match the metered demand restriction while using 10 MIT resulted in a loss of capacity. It is also interesting to note that the recovery of the throughput to match the demand again under time based metering was faster than under distance based metering in both peaks. (The plots were generated using Poisson arrival processes).

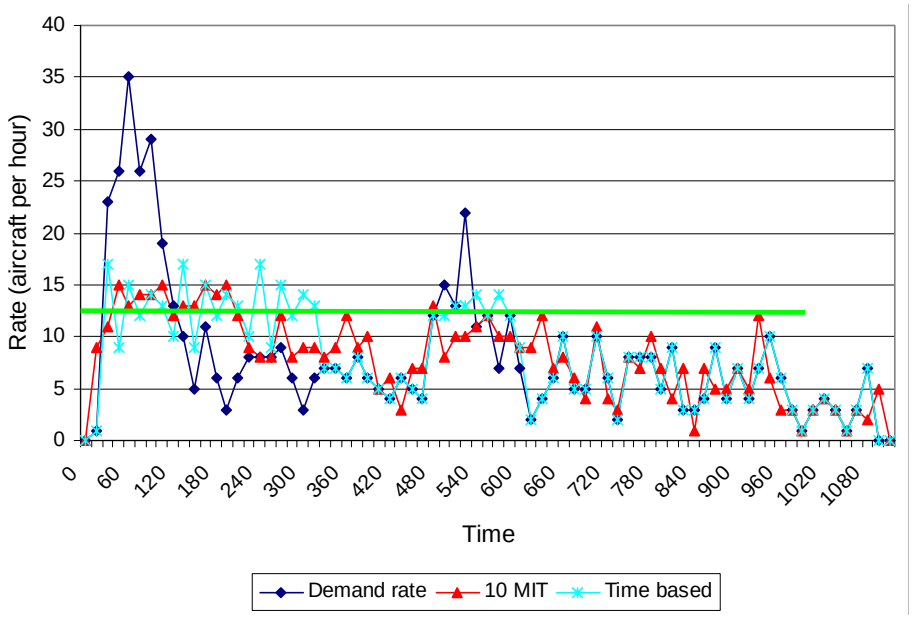

Figure 57. Better balancing between streams using time based metering

Modeling of High Resolution Metering

The high resolution metering is modeled explicitly by the resolution specified in the McTMA and baseline models, which is according to McTMA specifications (1 second but specified to the minute) and current operations (5 MIT), respectively.

6.3.3. Independence from Aircraft Speed Variability

In distance based metering, variability in aircraft speed leads to increased variability in the time of arrival at fixes because the restriction is specified in terms of distance separation. Figure 58 shows the effect of variability in speed on the delays (actual arrival time minus ETA) when MIT is applied to a flow of traffic with constant speed (across aircraft) and to a flow of traffic with variable speed (with an average value the same as the constant speed of the other flow). The flows were simulated using Poisson processes. It is evident as the figure shows that the variability in speed increases the delays. Time based metering however excludes this variability, as the metering assigns times of arrival at the fixes independent of speed.

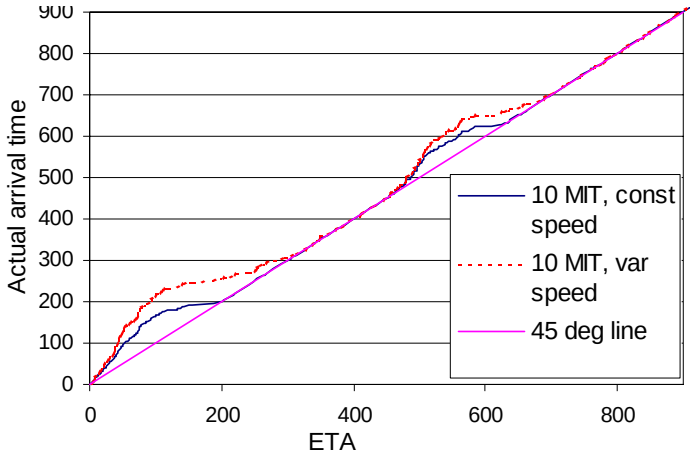

Figure 58. The effect of variability in speed on delays in distance based metering

The benefit mechanisms of time based metering independence from speed variability are presented in Figure 59.

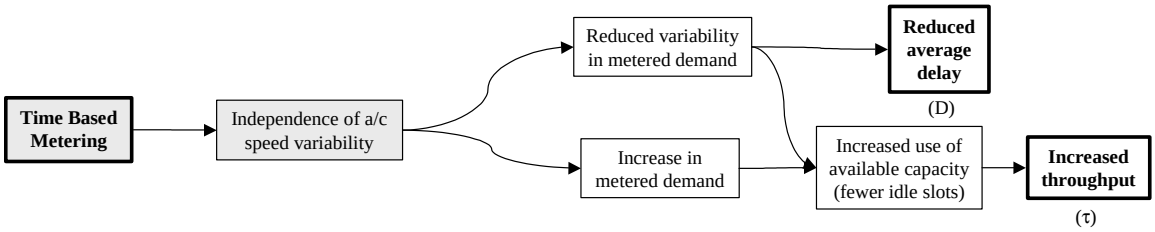

Figure 59. Benefit mechanism of independence of metering with speed variability

The independence of time based metering from aircraft speed variability increases the metered demand and reduces the variability in the metered demand. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the

constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Modeling of Metering Independence from Speed Variability

The modeling of this mechanism within time based metering is implicit to the calculation of STAs, discussed in Section 5.1. It is assumed that controllers will use appropriate advisories to meet the STA with a certain error. The effect of speed variability is captured in the baseline model by using a random distance separation between each pair of aircraft affected by a MIT restriction, and then using the actual aircraft speed to calculate its actual arrival time at a fix. The distance separation distribution for each MIT restriction magnitude is determined by the calibration of the MIT spacing baseline model as described in Section 4.1.2.

6.3.4. Freeze Horizon

The setting of a freeze horizon is a characteristic of time based metering, and refers to the specification of a time period from the current time, for each meter fix, within which the schedule is fixed for all aircraft (aircraft with STAs within that time period). The benefit mechanisms of setting the freeze horizon are presented in Figure 60.

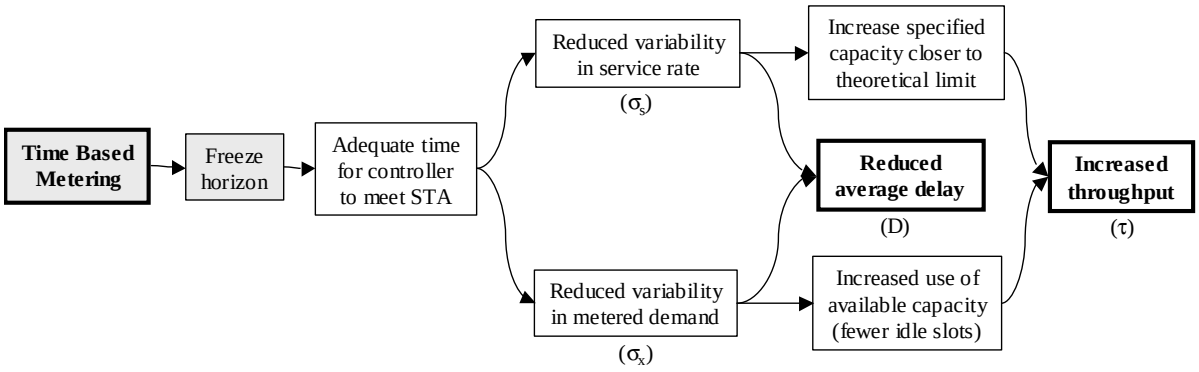

Figure 60. Benefit mechanism of setting the freeze horizon.

The longer the freeze horizon the more time the controllers have to meet STAs. However, the longer the freeze horizon the less time McTMA has to adjust STAs optimally according to changing constraints. There is therefore a tradeoff, and the optimal setting of the freeze horizon is a subject of research.

The freeze horizon ensures that the controllers working the aircraft scheduled at the meter fixes have enough time to absorb the delay required to meet the STA. Adequate time to meet the STA allows aircraft to be delivered to the fixes with high accuracy. Consequently, the variability in the service rate of each meter fix, and thus the variability in the metered demand for the next meter fix downstream is reduced.

A reduction in the variability in the service rate of a resource allows the specified capacity of that resource to be increased closer to the theoretical capacity. The theoretical capacity is a limit that should not be exceeded by the actual service rate in order to ensure safe operations. This is illustrated in Figure 61, where a reduction in the variability of

service rate (oscillating line) means that the specified capacity (the solid straight line) can be raised closer to the theoretical capacity (dashed straight line) without exceeding the safety limit. Applied to the constraining resource, this results in an increase in system throughput. System capacity increase can lead to a reduction in average delay, if demand is not increased or allows the demand to be increased to take advantage of the increased throughput. There is however a tradeoff between increased demand and reduced delay. This is discussed in detail in Section 8.

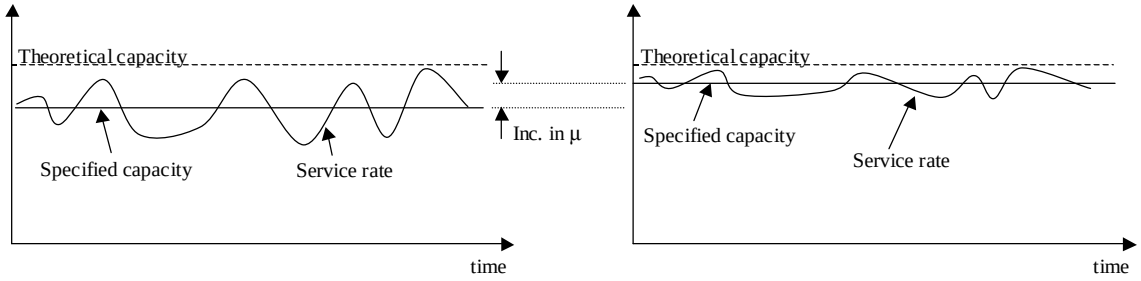

Figure 61. Effect of reduced variability in service rate

Reduced variability in service rate does also result in a reduction in delay directly however. This is illustrated in Figure 62. Variability can cause a low service rate followed by a high service rate, averaging to an average service rate over the entire period. If the demand were equal to this average service rate, those aircraft arriving when the service rate was low would incur delays as depicted in the upper part of the figure. If, however, the service rate were constant and equal to the average value throughout, none of the arriving aircraft would incur significant delays, as depicted in the lower part of the figure. The average delay is thus higher when the variability in service rate is high, even if the overall throughput is maintained at the same level.

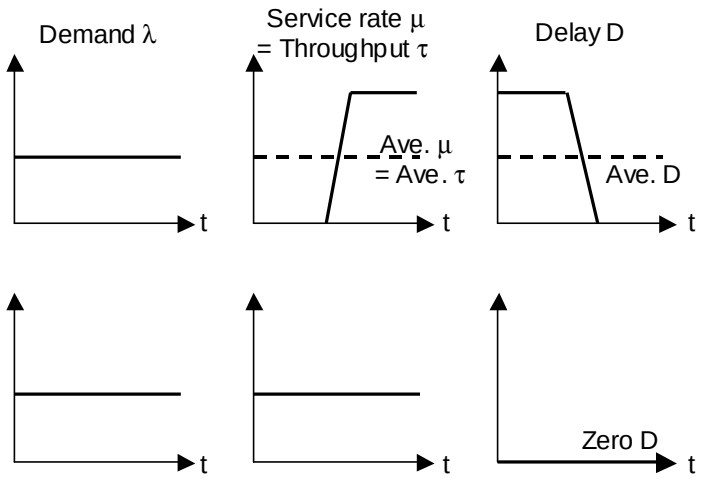

Figure 62. Higher delay resulting from high variability in service rate

The freeze horizon reduces the variability in the metered demand. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the

constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Modeling of Freeze Horizon

The error ε in meeting the STA in Equation (8) in Section 5.1 is a function of the length of the freeze horizon.

$$\varepsilon = f^{n}\left(freeze\ horizon\right) \tag{10}$$

The freeze horizons have been set for each meter fix following the procedure suggested by NASA's McTMA researchers – namely, aligning the freeze horizons with sector boundaries. The modeling of the freeze horizon is, therefore, implicit in the arrival flow network. The sector boundary alignment procedure is motivated by controller acceptability[8]. While it was intended to vary the freeze horizon as a parameter this procedure does not allow much variation. Variations of the freeze horizon may still be investigated in future research.

The error ε in Equation (10) has been set based on consultation with NASA's McTMA researchers and based on previous TMA field research to be a normal distribution centered around zero with a standard deviation of 90 seconds, as described in Section 5.1. This error was also varied in sensitivity analysis as described in Section 9.3.

6.4. Benefit Mechanisms of Delay Feedback and Capacity Distribution

Delay feedback and capacity distribution are functions of McTMA, and refer to the propagation of delays upstream according to the constraints on the capacities of the downstream airspace to absorb delay, and the generation of acceptance rate profiles according to STAs. The mapping of the benefits of delay feedback and capacity distribution is presented in Figure 63.

[8] Controllers prefer to command the aircraft the advisory needed to meet the STA as close to the handoff as possible, i.e. once the aircraft is under their control.

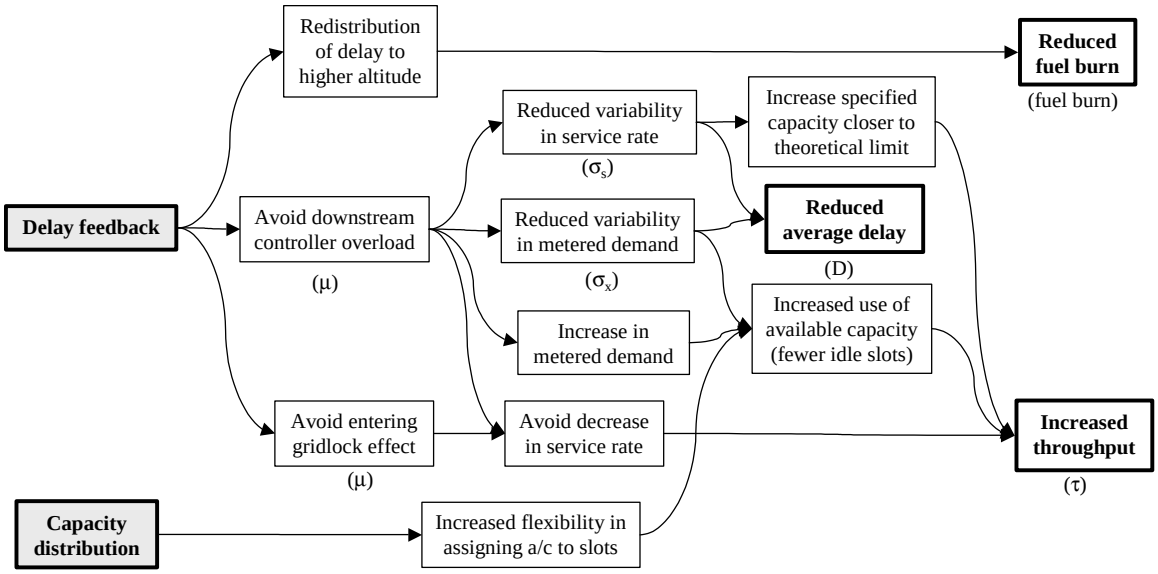

Figure 63. Benefit mechanisms of delay feedback and capacity distribution

Each mechanism in the Figure 63 is discussed in detail in the following subsections.

6.4.1. Redistribution of Delay to Higher Altitude

Because the function of delay feedback propagates delay upstream, and because aircraft descend during the arrival process, delay feedback results in the redistribution of delay to a higher altitude. The benefit mechanisms associated with redistribution of delays is presented in Figure 64.

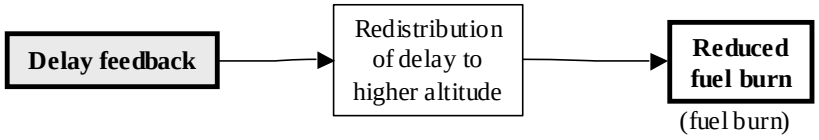

Figure 64. Benefit Mechanism of delay feedback: Redistribution of delay to a higher altitude

The redistribution of delay to a higher altitude results in reduced fuel burn, as aircraft burn more fuel at low altitude than high altitude.

Modeling of Redistribution of Delay to Higher Altitude

The reduction in fuel burn due to redistribution of delay to a higher altitude is modeled by calculating fuel burn resulting from delay, in each of the metering models. Fuel burn is a function of aircraft type, engine type, altitude and aircraft speed. For each aircraft and engine combination, aircraft manufacturers provide fuel burn tables, from which fuel burn can be calculated according to aircraft speed and altitude. Thus, in order to calculate fuel burn according to how much delay is absorbed for each meter fix, the following steps were followed:

- Associate an altitude and Mach number for a typical small, medium and large aircraft with each meter fix. The altitude and Mach number represent the altitude and Mach number at which delay would be absorbed for that meter fix. This was assumed to be the altitude and Mach number at the upstream meter fix.
- Using fuel burn model for the typical small, large, heavy and Boeing 757 aircraft, estimate the fuel burn rate, or fuel flow (FF, in lb/min) for each aircraft weight class, at each meter fix. Fuel burn rates for typical small, large, heavy and Boeing 757 aircraft were obtained from NASA CTAS Engineering Group.
- For each aircraft passing through a meter fix specify its fuel burn rate according to its weight class (small, large, heavy or Boeing 757 aircraft).
- Multiply the delay absorbed for each meter fix by the weighted average fuel burn rate FF_i for each aircraft, as illustrated in Figure 65 below. This yields the fuel burned during the absorption of delay for each aircraft at each meter fix.

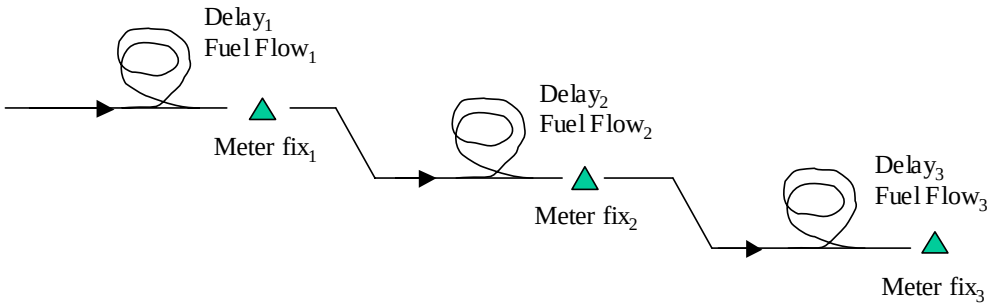

Figure 65. Fuel burn at meter fixes

6.4.2. Avoid Downstream Controller Overload

Delay feedback results in there being fewer requirements for delay absorption in the downstream sectors, as all the delay above the threshold values in each sector, is fed upstream. The delay threshold (delayability) is that amount of delay that can be absorbed without holding, so holding is greatly reduced through McTMA. This reduces workload on downstream controllers, and thus delay feedback avoids controller overload. Workload may be increased as a result for the upstream controllers; however the upstream airspace tends to have less capacity constraints and less congestion while the downstream airspace and resources such as the terminal area and the runways tend to be the flow bottleneck. The benefit mechanisms associated with avoiding downstream controller overload are presented in Figure 66.

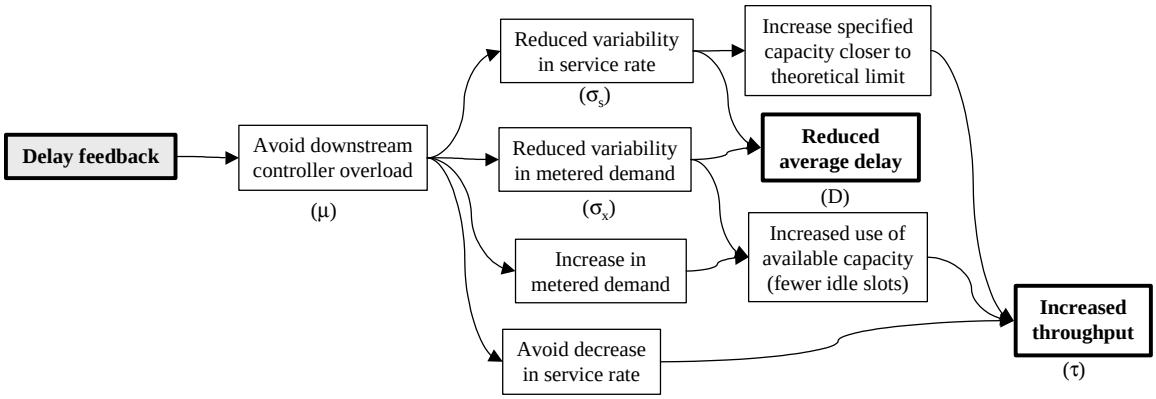

Figure 66. Benefit mechanisms of delay feedback: Avoiding downstream controller overload

When controllers become overloaded they are less effective in controlling aircraft and their service rate might drop as they miss slots because of being occupied by maintaining separations and holding aircraft. A reduction in the service rate of the constraining resource results in a decrease in throughput. Thus, by avoiding downstream controller overload a decrease in service rate is avoided and average throughput is increased. Similarly, by avoiding downstream controller overload fewer gaps are opened in the arrival stream and the metered demand increases, leading to an increase in the throughput.

When controllers become overloaded the accuracy with which aircraft can be metered is also reduced. Consequently, by avoiding downstream controller overload McTMA enables reduced variability in the service rate of the resource where metering is applied, as well as a reduction in the variability of the metered demand for the downstream resources. As discussed in Section 6.3.4 such a reduction in service rate variability allows for an increase in the specified capacity closer to the theoretical limit of capacity, which results in increased throughput if applied to the constraining resource. Reductions in the variability of the service rate, and of the metered demand, both result in a reduction in the average delay, directly, also as discussed in Sections 6.3.1 and 6.3.4. Reduced variability in demand also results in an increase in the use of available capacity as discussed in Sections 6.3.1. Applied to the constraining resource this results in fewer idle slots and thus increased throughput.

Modeling of Avoidance of Downstream Controller Overload

The effect of controller overload is captured using the capacity model presented in Section 3.1, where the throughput was shown to drop as the congestion and delay levels become excessively high because of, among other factors, high controller workload. Using McTMA's delay feedback mechanism, the congestion and delay levels will be maintained below the levels at which the throughput drops. In addition because of lower workload it is expected that the capacity constraints imposed in McTMA can be set closer to the theoretical safety limit. This is discussed in more detail in Section 6.3.

6.4.3. Avoid Gridlock Effect

Delay feedback results in there being less congestion in the downstream sectors, and thus enables gridlock to be avoided. Gridlock refers to when no aircraft are able to advance in the flow because of interdependence between streams (usually arrival and departure streams) each waiting for the other to advance at the same time. Gridlock rarely happens in reality because the system responds before it is able to completely manifest. The benefit mechanism associated with avoiding the gridlock effect is presented in Figure 67.

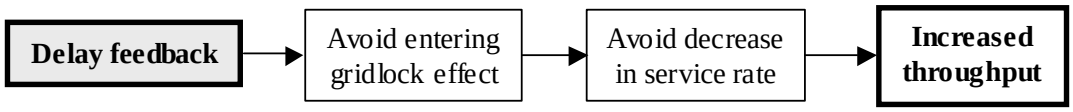

Figure 67. Benefit mechanism of delay feedback: Avoiding gridlock effect

Even though gridlock rarely materializes, as the system approaches gridlock the service rate drops in an attempt to avert it. Thus, avoiding approaching the gridlock effect results in avoiding decreasing the service rate and thus allowing for an increased system throughput.

Modeling of Avoidance of the Gridlock Effect

The gridlock effect is also be captured in the capacity model presented in Section 3.1, where the throughput is expected to drop as the congestion and delay levels become excessively high because of, among other factors, approaching gridlock. Using McTMA's delay feedback mechanism, the congestion and delay levels are maintained below the levels at which the throughput drops.

6.4.4. Increased Flexibility resulting from Capacity Distribution

Capacity distribution and the use of the rate profiler to calculate acceptance rate profiles according to STAs propagated from downstream allows for the specification of new STAs according to local constraints, ETAs, and the acceptance rate profile. This allows for greater flexibility in assigning aircraft to slots, as aircraft must meet local constraints and the specified acceptance rate profile only, and not the sequencing requirements propagated from downstream. The benefit mechanism associated with increased flexibility resulting from capacity distribution is presented in Figure 68.

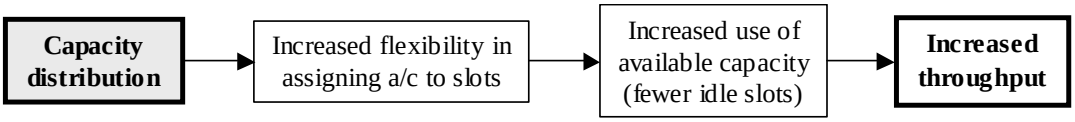

Figure 68. Benefit mechanisms of capacity distribution: Increased flexibility

Increased flexibility in assigning aircraft to slots results in increased use of the available slots, as there are fewer constraints to satisfy in filling slots. This results directly in increased throughput.

Modeling of Increased Flexibility

The increased flexibility in assigning aircraft to slots is captured explicitly in the implementation of capacity distribution in the McTMA model as described in Section 5.1.

6.5. Benefit Mechanisms of Dynamic Metering

Dynamic metering is a function of McTMA and refers to the dynamic adjustment of STAs. McTMA regenerates STAs every 12 sec, up to the freeze horizon (a set distance from the meter fix that typically equates 19 minutes) after which the STAs become fixed. This is in contrast to current operations where a MIT restriction generally remains unchanged for up to few hours, as detailed in Section 3.1.2. The benefit mechanisms of dynamic metering are presented in Figure 69.

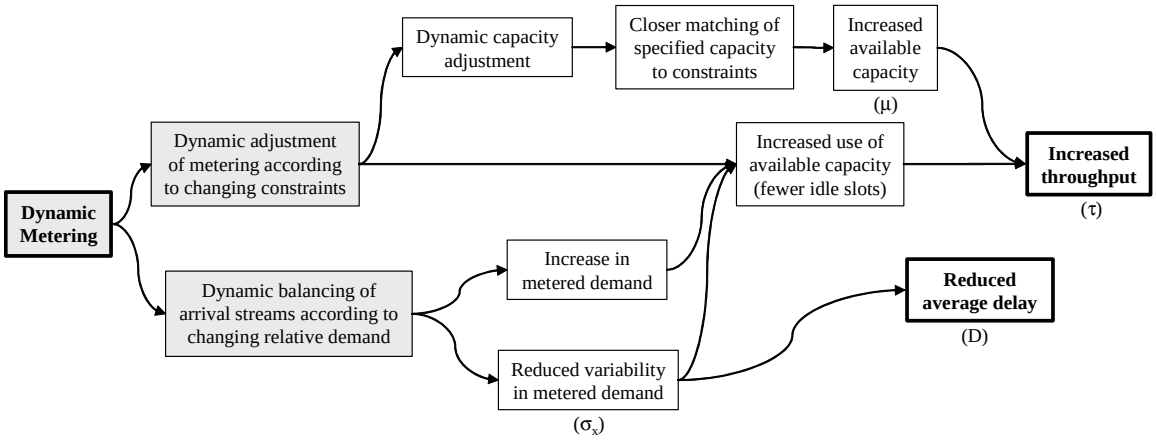

Figure 69. Benefit mechanisms of dynamic metering

Dynamic metering includes the dynamic adjustment of metering according to changing constraints, and dynamic balancing of arrival streams according to changing relative demand. The benefit mechanisms associated with each of these characteristics are presented on the following subsections.

6.5.1. Dynamic Adjustment of Metering According to Changing Constraints

Dynamic adjustment of metering according to changing constraints is a characteristic of dynamic metering. Constraints do change, and current operations are not well equipped to adjust the metering accordingly. This is because of lack of automation and the slow system response to changes in the distance based MIT restrictions. McTMA's dynamic metering however allows almost immediate changes in the metering to accommodate any changing constraints. The benefit mechanisms associated with dynamic adjustment of metering according to changing constraints is presented in Figure 70.

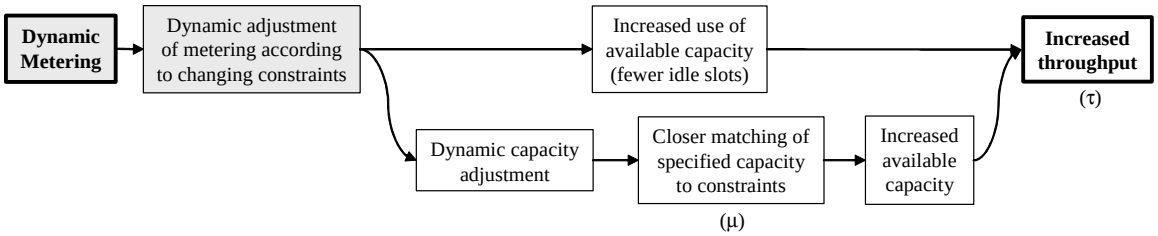

Figure 70. Benefit mechanism of dynamic metering: Dynamic adjustment of metering according to changing constraints

When constraints become more restrictive, they must be accommodated for safety reasons. Thus, a decrease in capacity does result in a corresponding increase in metering, even in current operations. Because restrictions cannot be changed often and rapidly, the flow is generally restricted according to the most severe level of the constraints expected. This means that capacity is lost in those periods where the constraints are not as severe. Dynamic metering allows more immediate response to changes in capacity constraints. Use of the available capacity of a resource, and particularly that of the constraining resource, is thus increased. This reduces the number of idle slots in the stream, increasing throughput.

Dynamic adjustment of metering according to changing constraints allows for dynamic adjustment of the specified capacity, since the metering can respond, which is currently not the case. Specified capacity could be changed dynamically according to dynamic changes in the actual capacity, allowing close matching of the capacity to the constraints. In the case of the runway, the capacity envelope could be utilized to a greater extent, adjusting in McTMA specified AARs according to the arrival and departure schedules. This would allow an increase in available capacity, which, if applicable to the constraining resource, would allow for an increase in throughput.

Modeling of Dynamic Adjustment of Metering According to Changing Constraints

The dynamic adjustment of the metering according to changing constraints is modeled through matching variable acceptance rates. Acceptance rates vary over the day as the reported capacity is changed due to changes in runway configuration or weather. In each configuration the acceptance rate is also varied as a function of departure rate according to the capacity envelopes described in Section 3.1.1. McTMA metering is applied over periods of 30 minutes or longer as detailed later in Section 7.2. On the other hand, in the baseline model the MIT restrictions are maintained for a period on the order of hours (at least 1 hour as observed from analyzing the facility restriction logs).

While STAs were not updated every 12 seconds, the dynamic nature of McTMA's time based metering was captured by updating the STAs between upstream tiers and downstream tiers adjusting for accumulated errors in meeting the upstream STAs. The error in meeting the STA was added to each STA, at each tier as described in Section 5.1. However, because McTMA updates ETAs and recalculates STAs every tier, it is able to adjust for the errors incurred in upstream tiers. This was modeled by sampling an error from the error distribution, adding it to the STA of the tier in which the error was incurred, and then subtracting it from the amount of delay to be absorbed in the tiers downstream. The change in the amount of delay to be absorbed in each downstream tier

was limited in that the resulting adjusted amount of delay to be absorbed must be greater than zero, and less than the delayability of the tier.

6.5.2. Dynamic Balancing of Arrival Streams

Dynamic balancing of arrival streams according to changing relative demand is a characteristic of dynamic metering. This refers to the accurate balancing of flow from the different streams as the demand on each stream changes. Even though current operations attempt to place lower restrictions on the heavy demand flow when applying MIT restrictions, they are not dynamic enough to match changing relative demand. The benefit mechanisms associated with dynamic balancing of arrival streams are presented in Figure 71.

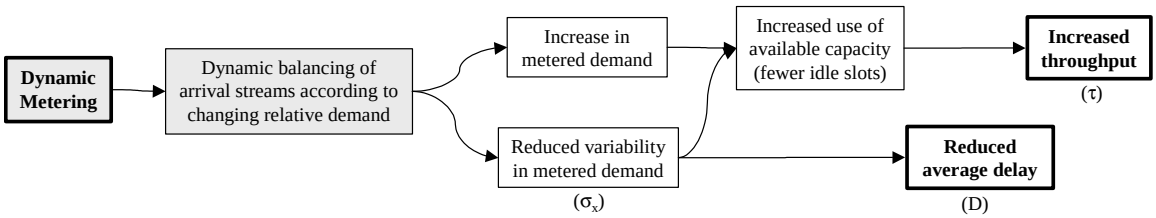

Figure 71. Benefit mechanisms of dynamic metering: Dynamic balancing of arrival streams

Because McTMA is able to dynamically balance between flows as demand changes, the metered demand is increased and its variability is reduced. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Modeling of Dynamic Balancing of Arrival Streams

The dynamic balancing of arrival streams is modeled implicitly through matching variable arrival rates from different streams and through the coordination between multiple streams, with all aircraft on all streams accounted for in the calculation of STAs. In the baseline model, heavy demand streams are assigned lower MIT restrictions as described in Section 4.1.1.

6.6. Benefit Mechanisms of Tiered Metering

Tiered metering is a function of McTMA, and refers to metering at a number of meter fixes on each aircraft stream, extending a number of tiers upstream from the destination airport. The benefit mechanisms of tiered metering are presented in Figure 72.

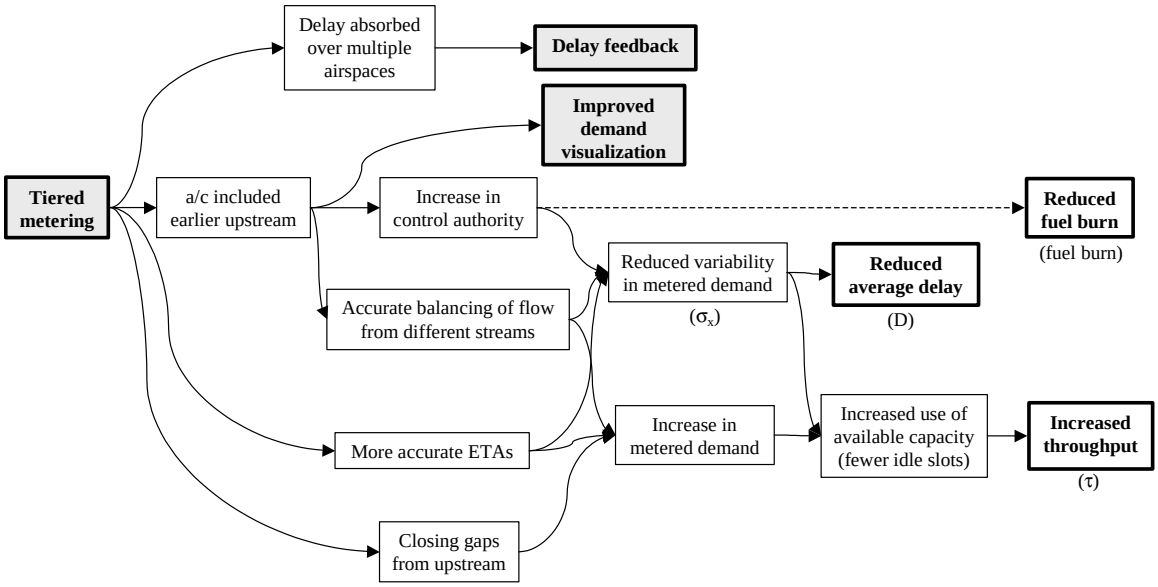

Figure 72. Benefit mechanisms of tiered metering

Each mechanism in Figure 72 is discussed in detail on the following subsections.

6.6.1. Delay Absorbed over Multiple Airspaces

Tiered metering allows for the absorption of the delay to be specified over a number of airspaces, and not restricted to the airspace in or near the TRACON boundary. The benefit mechanism associated with the absorption of delay over multiple airspaces is presented in Figure 73.

Figure 73. Benefit mechanism of tiered metering: Delay absorption over multiple airspaces

The absorption of delay over multiple airspaces enables the McTMA function of delay feedback, as delay can be fed back through multiple meter fixes, ensuring that the delay threshold within each airspace is not violated. The benefits and modeling of delay feedback are discussed in detail in Section 6.4.

6.6.2. Inclusion of Aircraft Earlier Upstream

Multiple tiers of metering allow aircraft to be included in the calculation of ETAs and STAs earlier upstream, and allows more aircraft to be included in the metering process. The benefit mechanisms associated with this is presented in Figure 74.

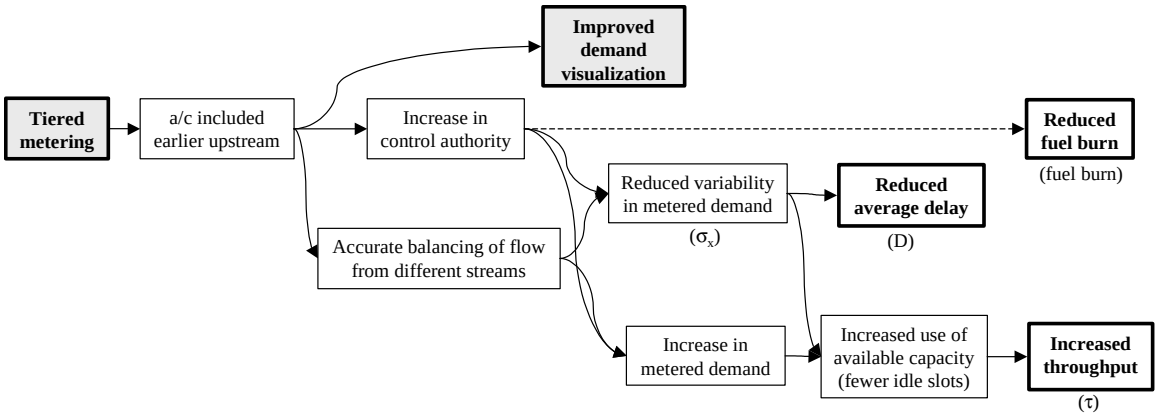

Figure 74. Benefit mechanism of tiered metering: Inclusion of aircraft earlier upstream

The inclusion of more aircraft upstream and the inclusion of each aircraft earlier in the arrival process improves visualization of demand. Demand visualization is a function of McTMA, and is discussed in detail in Section 6.7. This discussion includes a discussion on the function's benefits mechanisms and modeling of these mechanisms.

Inclusion of aircraft earlier upstream in the arrival process also increases the control authority which McTMA has on the aircraft. This is because ETAs and STAs can be generated earlier. It is hypothesized that increased control authority may reduce fuel burn, as the longer an aircraft is under McTMA control authority, the more distance the controller has to absorb the delay that ultimately needs to be absorbed. The delay can thus not only be absorbed at higher altitude, as discussed in Section 6.4.1, but also through more fuel efficient delay absorption strategies, such as speed reduction as opposed to vectoring or holding. This benefit has yet to be confirmed however and is thus currently not to be modeled.

The inclusion of aircraft earlier upstream in the arrival process also allows for more accurate balancing of the flow from different streams. This, along with increased control authority, increases the metered demand and reduces its variability. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Modeling of Inclusion of Aircraft Earlier Upstream

The modeling of tiered metering, and the inclusion of aircraft earlier upstream is implicit to the modeling of McTMA and time based metering described in Section 5.1, and ensured by the identification of the meter fixes in Appendix A.

6.6.3. Closing Gaps from Upstream

Tiered metering also allows for gaps in the flow from upstream to be closed by rescheduling a new set of STAs at the next meter fix downstream. Thus, if a flight is unable to meet its STA, the resulting gap can be closed at the next meter fix. Without tiered metering, and thus only one meter fix, any slots scheduled at the meter fix, but

ultimately missed, cannot be recovered. The benefit mechanism associated with this is presented in Figure 75.

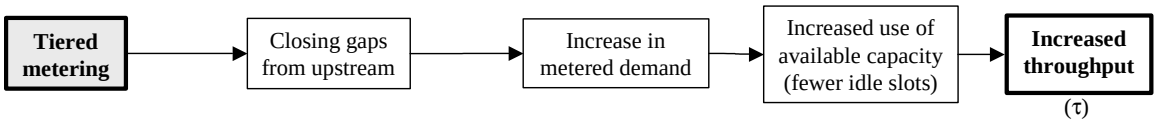

Figure 75. Benefit mechanism of tired metering: Closing gaps from upstream

Closing of gaps in the flow from upstream increases metered demand. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput.

Modeling of Closing Gaps from Upstream

The modeling of the closing of the gaps propagating down from upstream is implicit to the modeling of time based metering and the generation of STAs according to minimum delay at each meter fix. It is also captured in the adjustment of STAs between tiers as described in modeling dynamic metering.

6.6.4. More Accurate ETAs

Tiered metering allows for more accurate ETAs, as the distance between the system outmost boundary and the destination runway is divided into smaller tiers. The smaller tiers allow more accurate prediction of ETAs for each tier and more feasible STAs to be computed based on the more accurate ETAs. The benefit mechanism associated with this is presented in Figure 76.

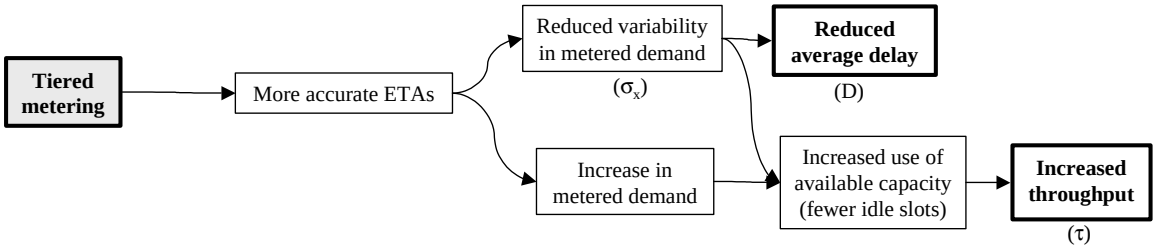

Figure 76. Benefit mechanism of tired metering: More accurate ETAs

More accurate ETAs ultimately result in an increase in the metered demand and a reduction in its variability as STAs specified according to the more accurate ETAs are more feasible, and can thus be met more accurately. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Modeling of More Accurate ETAs

The modeling of more accurate ETAs resulting from tiered metering is implicit in modeling unimpeded transition times for each tier as opposed to a single transition time

for the whole system, as described in Section 3.2, thus reducing the variability in ETAs. This variability is further reduced by taking into account flight distance, wind and aircraft class.

6.7. Benefit Mechanisms of Demand Visualization

Demand visualization is a function of McTMA, and refers to the visualization of demand through loadgraphs and timelines, showing demand according to ETAs and STAs, in addition to displaying the capacity and expected delay. The benefit mechanisms of demand visualization are presented in Figure 77.

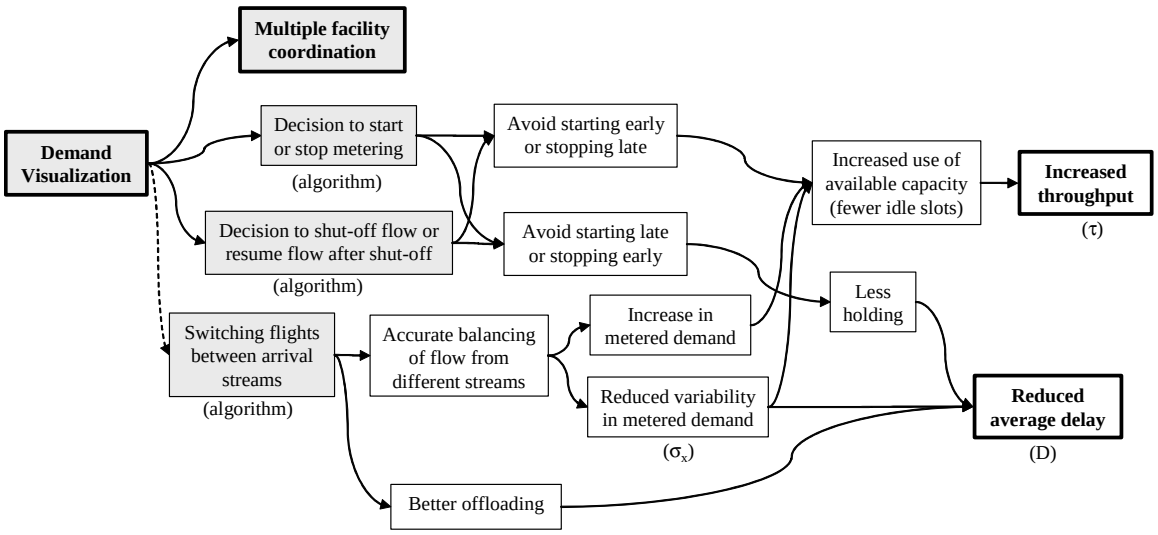

Figure 77. Benefit mechanisms of demand visualization

Because each facility is able to visualize demand from a number of other facilities a more global view is created and coordination between the facilities is improved. The benefits of multiple facility coordination are discussed in detail in Section 6.8.

Demand visualization includes the following characteristics:

- 'What if' functionality.
- Decision making regarding starting and stopping metering, and shutting off and resuming flow.
- Switching flights between arrival streams.

The benefit mechanisms associated with each of these characteristics are presented on the following subsections.

6.7.1. What-if Functionality

The 'what-if' functionality of McTMA refers to the potential use of McTMA to test MIT restrictions, before a final MIT restriction is implemented. By inputting the MIT restriction to be tested, McTMA provides an indication of resulting delay. This use of McTMA is not currently planned but would represent the potential benefits during a transition phase between current operations and time based metering operations. In this

transition phase only the visualization tools of McTMA would be deployed and made available to traffic managers. Metering would thus still be according to MIT.

It was intended to analyze the benefits of demand visualization in this transition deployment phase through modeling a McTMA what-if capability. Modeling this capability would also provide an estimate of the potential benefits if such a function were added to McTMA at a later time. However, due to time constraints and because such a transition phase is no longer planned, NASA McTMA researchers advised not to perform this analysis in favor of other more relevant extensions. The benefit mechanisms derived are kept in this description for completeness and future reference.

The benefit mechanism of the 'what if' functionality is presented in Figure 78.

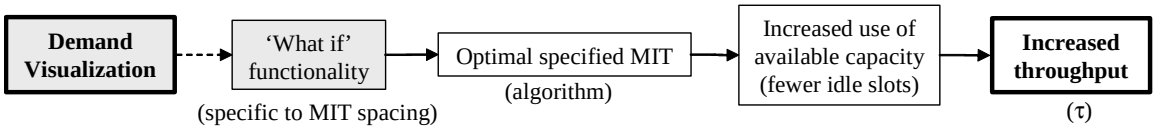

Figure 78. Benefit mechanism of 'what-if' functionality

By comparing the resulting delay from a number of different restrictions, McTMA can enable the optimal MIT restriction to be applied. Note that the MIT restrictions must still be specified in increments of 5 MIT. Optimization of the MIT restrictions would apply a rate that better matches the constraints, and hence increases the use of the available capacity by reducing the number of idle slots. Applied to the constrained resource, this will lead to increased throughput.

Modeling of 'What if' Functionality

The 'what if' functionality of McTMA was not modeled in this benefit assessment, but may be modeled in the future through the specification of optimal MIT restrictions. This requires the comparison of the results of a number of different MIT restrictions. Given the demand and other constraints, the delay resulting from a number of different MIT restrictions surrounding the originally specified MIT restriction would thus be calculated and compared. The restriction resulting in the lowest delay would then be selected for application.

6.7.2. Decisions Regarding Starting and Stopping Flows and Metering

Because demand visualization enables more accurate demand in the future to be seen by traffic managers, decision making on when to start and stop metering, and decision making on when to shut-off flow and when to resume flow after a shut-off, can be improved. The benefit mechanisms of decision making regarding the starting and stopping of flows and metering is presented in Figure 79.

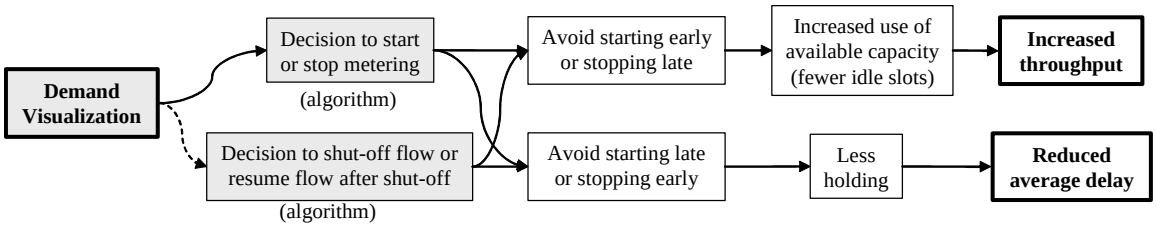

Figure 79. Benefit mechanism of decision making regarding starting and stopping flows and metering

Demand visualization will allow for better timed starting and stopping of metering. Avoidance of starting metering too early or stopping too late will result in an increase in the use of the capacity of the resource, as the resource's capacity will be utilized for as long as it is available. Applied to the constraining resource, this will result in an increase in throughput.

Avoiding stopping metering too early or starting too late will result in less requirement for holding (or other means of managing the excess demand), which results in reduced average delays.

Modeling of Decision Making Regarding Starting and Stopping Flows and Metering

Only starting and stopping metering, both time based and distance based, were modeled. Based on analysis of facility logs shut-off is employed in extreme conditions in transition to more restrictive flow management programs such as a ground delay program. Most of the analysis was conducted on normal conditions and moderately restrictive conditions when only MIT was applied. Extension to more severe and restrictive conditions was intended but time did not permit and may be considered for future extension.

Based on consultation with NASA's McTMA researchers, time based metering was started when demand was predicted to exceed the reported capacity of the runway configuration as detailed in Section 7.2. Predicted demand was modeled using the ETAs as specified in Section 3.2 and an algorithm was developed to detect when predicted demand exceeds capacity. Also based on consultation with NASA's McTMA researchers time based metering was stopped when the delay dropped to zero. A minimum duration of 30 minutes was applied. In the baseline model distance based MIT restrictions were applied during their actual time of application according to the facilities restriction logs. Also based on consultation with NASA's McTMA researchers a scenario was run where time based metering replaced MIT during the same periods when MIT was applied, as detailed in Section 7.2 also.

6.7.3. Switching Flights between Arrival Streams

McTMA enables visualization of demand on all arrival streams, and thus, decisions on switching flights from one arrival stream to another can be made more effectively. This should improve the effectiveness of offloading aircraft from ZOB heading to ZNY, through ZBW or ZDC as currently practiced. It should also encourage more rerouting through other facilities which is currently avoided because of difficulty in

coordination. This is particularly useful at times when certain gates are affected by bad weather and other gates are not. The benefit mechanism of switching flights between arrival streams is presented in Figure 80.

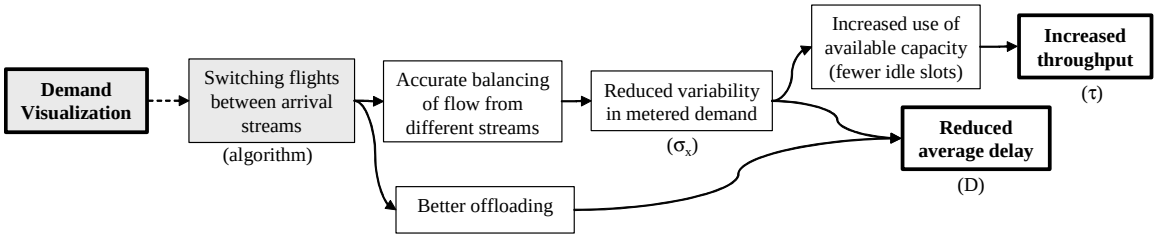

Figure 80. Benefit mechanism of switching flights between arrival streams

Switching flights between arrival streams results in more accurate balancing of flow from different streams, which results in increased metered demand and reduction in metered demand variability. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1. Switching flights between arrival streams also results in better and more optimal offloading which results in less delay as shorter reroutes are used.

Modeling of Switching Flights Between Arrival Streams

Switching flights between streams was not modeled based on feedback from NASA's McTMA researchers that rerouting of flights between arrival gates is not currently considered as a McTMA function. The extension to model and analyze such a function may be considered in future extension work. Particularly, since traffic managers indicated during the site visits that it would be beneficial if the tool advised rerouting options when certain gates are affected by severe weather, and the facility logs contained indication that airlines or pilots get dismayed because of offloading through longer routes resulting in more fuel burn.

6.8. Benefit Mechanisms of Multiple Facility Coordination

Multiple facility coordination is a function of McTMA, and refers to the coordination between all facilities involved in the arrival process that is enabled by McTMA. The benefit mechanisms of multiple facility coordination are presented in Figure 81. Multiple facility coordination includes coordination between the TRACON and the associated ARTCCs, coordination between upstream and downstream facilities, coordination between facilities at the same tier, and coordination with facilities outside the McTMA system such as the Command Center (ATCSCC) and other ARTCCs. The benefit mechanisms of each of these features are presented in the following subsections.

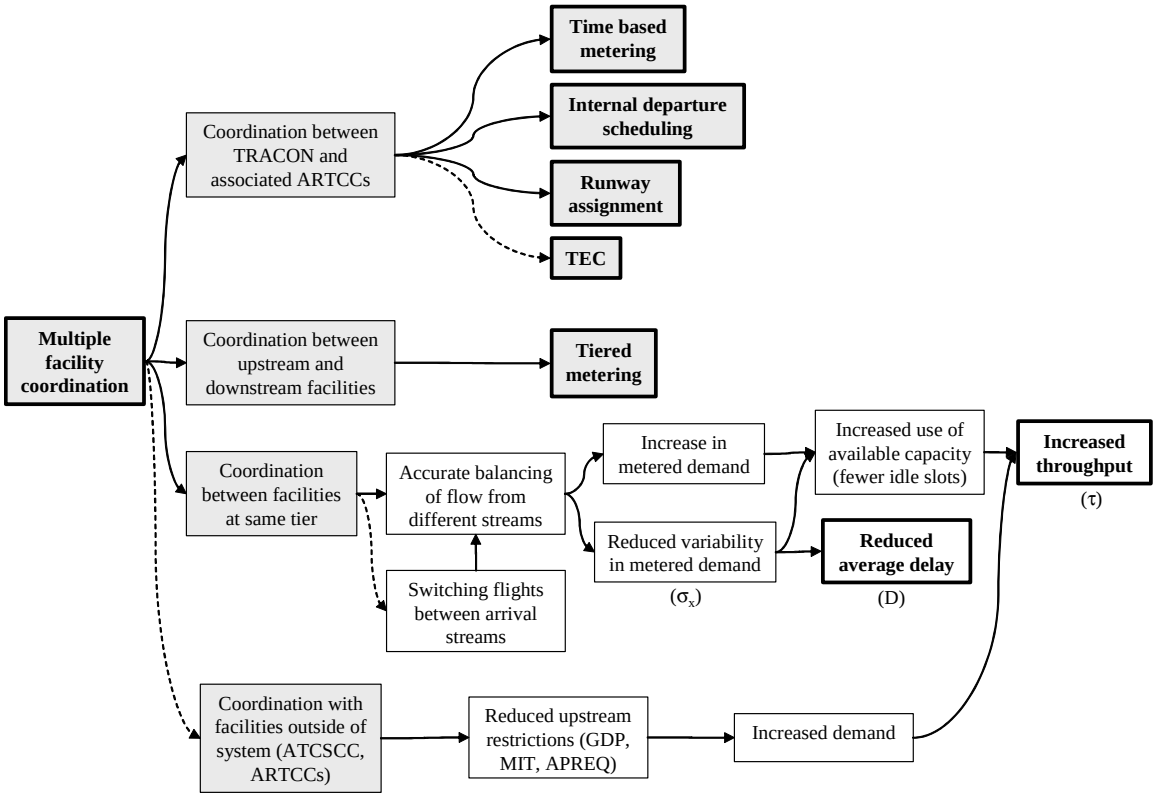

Figure 81. Benefit mechanisms of multiple facility coordination

6.8.1. Coordination between TRACON and Associated ARTCCs

The benefit mechanisms of coordination between the TRACON and associated ARTCCs, which in this case are ZNY, ZBW, ZOB and ZDC, are presented in Figure 82.

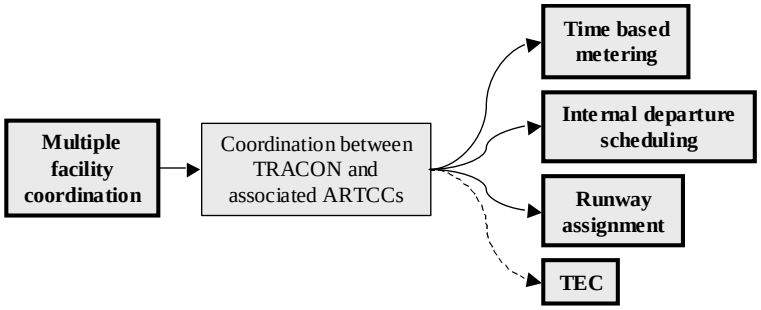

Figure 82. Benefit mechanisms of coordination between TRACON and associated ARTCCs

A number of functions of McTMA are enabled by the coordination between the TRACON and ARTCCs involved in the arrival process. Particularly this includes time based metering (including delay feedback and capacity distribution), internal departure scheduling, runway assignment; and tower en-route control. The benefits and modeling of time based metering is discussed in detail in Section 5.1 above, while the benefits of the runway assignment and internal departure scheduling functions are discussed in detail in Sections 6.9 and 6.10 below. McTMA gets information about TEC through ETMS

because they are not tracked by the Host computer. To accommodate for these aircraft, the acceptance rates and other constraints are adjusted according to the number of TEC aircraft expected.

Modeling of Coordination between TRACON and Associated ARTCCs

Coordination between the TRACON and Centers is likely to improve operations for each of the functions described. Because these functions are enabled by the coordination function, however, the modeling of the benefits of coordination is implicit to the modeling of the functions.

Based on consultation with NASA's McTMA researchers, in the McTMA model TEC arrivals are assigned ETAs according to their actual time of arrival at the runway. Then they are delayed if needed by time based metering. In this manner TEC traffic is considered as part of the constraint reserving landing slots. This is an approximation of the effect of TEC traffic since the McTMA approach to them is a subject of research.

6.8.2. Coordination between Upstream and Downstream Facilities

The benefit mechanism of coordination between upstream and downstream facilities is presented in Figure 83.

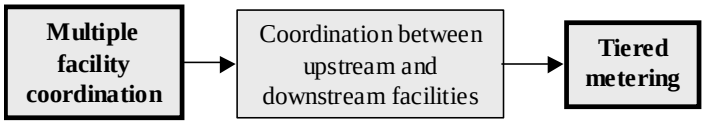

Figure 83. Benefit mechanism of coordination between upstream and downstream facilities

Coordination between upstream and downstream facilities enables tiered metering. Without this coordination metering over a number of tiers would not be possible.

Modeling of Coordination between Upstream and Downstream Facilities

In the modeling of tiered metering, coordination is thus implicit.

6.8.3. Coordination between Facilities at the Same Tier

The benefit mechanism of coordination between facilities at the same tier is presented in Figure 84.

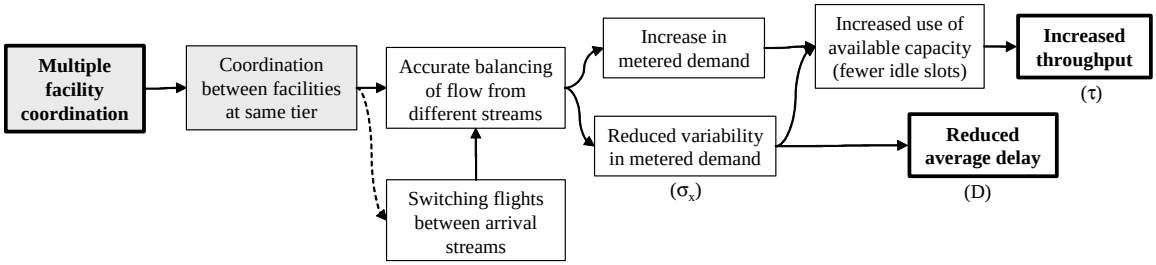

Figure 84. Benefit mechanism of coordination between facilities at the same tier

Coordination between facilities at the same tier enables coordination of the scheduling between different facilities. This results particularly in more accurate balancing of the flow from different streams and switching of flights between arrival streams, which increase the metered demand and reduce its variability. As discussed in Section 6.3.1 this results in an increase in the use of available capacity. Applied to the constraining resource this results in fewer idle slots and thus increased throughput. Reduced variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

Modeling of Coordination between Facilities at the Same Tier

Coordinated scheduling is again implicit to the generation of the schedule of STAs according to the demand on all the arrival streams, and to the balancing of flows from different streams where delay is propagated back through merging streams, as described in Section 5.1. Switching of flights between stream is not modeled as described under demand visualization.

6.8.4. Coordination with Facilities outside the System

McTMA's tiered time based metering may have an impact on the NAS outside the McTMA system. This impact materializes through coordination and interaction between the McTMA facilities and facilities outside the McTMA system. For example, through interaction with the Command Center, traffic flow management programs such as the Ground Delay Program (GDP) may be reduced in number and severity because of the application of McTMA. Through interaction and coordination with ARTCCs upstream of the McTMA system MIT restrictions and corresponding APREQ restrictions may be reduced as well because of the application of McTMA. These interactions result in the benefit mechanisms represented in Figure 85.

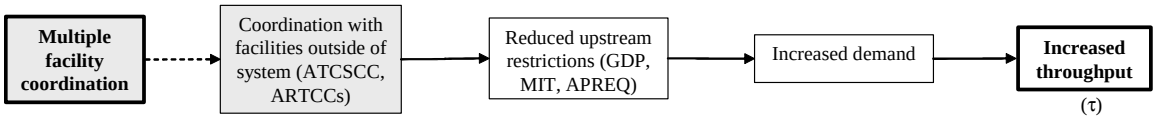

Figure 85. Benefit mechanism of coordination with facilities outside the system

Reduction in the GDP and other restrictions on traffic outside the McTMA system increases the demand on the McTMA system. This is because the arrival traffic enters the McTMA system less restricted by ground delay and MIT. The increase in demand for the McTMA system takes advantage of the increased capacity resulting from time based metering and reduced downstream controller workload. This results in an increase in the system throughput.

Modeling of Coordination with Facilities outside the McTMA System

Based on consultation with NASA's McTMA researchers it was decided, conservatively and due to time and resource constraints, not to investigate and analyze the impact of McTMA on the other restriction programs, such as GDP, applied in the NAS outside of the McTMA system. This could be done for example by increasing the

demand for the McTMA system to take advantage of the increased capacity due to time based metering. Therefore, the arrival and departure demand for the McTMA system was kept unchanged from the current actual level of traffic, as detailed in Section 7.2. However, the benefits assessment was conducted in periods of time when neither GDP for the analyzed airports nor MIT restrictions on the boundary of the system, were applied.

6.9. Benefit Mechanisms of Runway Assignment

Runway assignment is a function of McTMA, and refers to the assignment of runways to arriving flights in such as way as to reduce delay in the entire system. This function is not currently implemented for McTMA. Therefore, its analysis was not conducted. The benefit mechanisms of runway assignment are presented in Figure 86 for completeness and to support future extension.

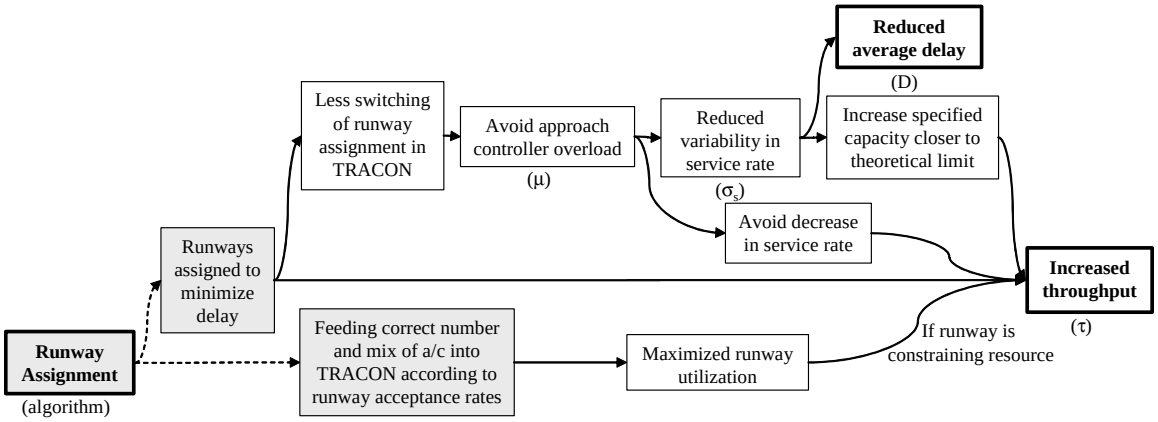

Figure 86. Benefit mechanisms of runway assignment

The runway assignment function of McTMA has two primary characteristics. Runways are assigned to minimize delay, and the correct number and mix of aircraft are fed into the TRACON according to specific runway acceptance rates. These characteristics are discussed in detail in the following subsections.

6.9.1. Landing Aircraft on Optimal Runways to Minimize Delay

The benefit mechanism of landing aircraft on the optimal runways to minimize delay is presented in Figure 87.

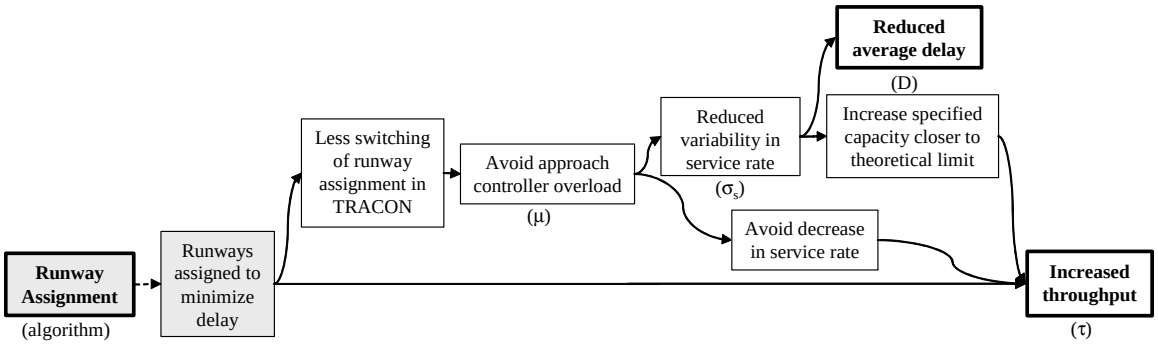

Figure 87. Benefit mechanism of landing aircraft on the optimal runways to minimize delay

Minimization of delay in runway assignment results directly in an increase in throughput as landings are packed closer together closing gaps. However, the landing of arrivals on the optimal runways also results in less switching of the runway assignment in the TRACON, as the runway that is assigned is optimal. Less switching of the runway assignment results in improved approach controller operations, and controller overload is thus avoided. This means that the reduction in service rate associated with controller overload, and the corresponding reduction in throughput, is avoided. Improved controller performance also means that aircraft are served more consistently, and the variability in service rate is thus reduced. As discussed in Section 6.3.4 this leads to a direct reduction in average delay, and also allows for an increase in the specified capacity closer to the theoretical limit, which would increase throughput.

Modeling of Landing Aircraft on Optimal Runways to Minimize Delay

The assignment of runways is not modeled in this benefits assessment, but may be modeled in the future explicitly according to the runway assignment algorithm in [4]. This algorithm assigns a runway to minimize the delay of all the aircraft entering the TRACON.

Aircraft should initially be assigned runways according to their type, their arrival fix, the airport runway configuration, and the acceptance rates on each runway. If runway configuration should change, all aircraft in the air may be re-assigned default runways according to the new configuration. If the acceptance rate on any runway should change, the aircraft flying between the first and second tiers of meter fixes (i.e. the next fix that they will cross is their arrival fix) may be rescheduled.

6.9.2. Feeding correct number and mix of Aircraft into TRACON

The benefit mechanism of feeding the correct number and mix of aircraft into the TRACON is presented in Figure 88.

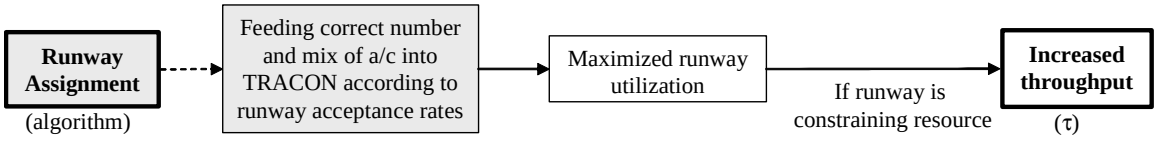

Figure 88. Benefit mechanism of feeding the correct number and mix of aircraft into the TRACON

Runway assignment results in the correct number of aircraft and the correct mix of aircraft, with respect to jets and turboprops, being fed into the TRACON in order to maintain pressure on the runways. This results in maximized runway utilization, which, if the runway is the constraining resource, results in increased throughput. In addition, in current operations the lack of information about runway assignment prevents the traffic managers from favoring aircraft heading towards non-restricted runways. With the McTMA runway assignment and the knowledge about which runway the aircraft is landing on, aircraft not affected by a restriction can be allowed to flow freely increasing utilization of the runways.

Modeling of Feeding Correct number and mix of Aircraft into TRACON

As detailed above, in Section 6.9.1, runway assignment may be explicitly modeled according to the runway assignment algorithm in [4].

6.10. Benefit Mechanisms of Internal Departure Scheduling

The benefit mechanisms of McTMA internal departures release are presented in Figure 89.

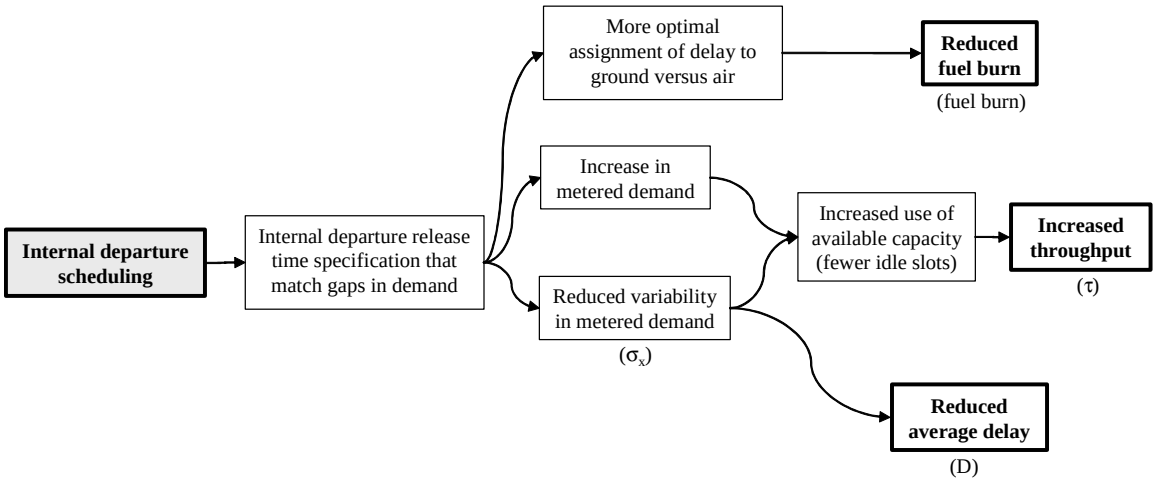

Figure 89. Benefits mechanism of internal departure scheduling

The scheduling of internal departures in McTMA results in the specification of internal departure release times that accurately match gaps in the arrival stream. This increases the metered demand and the use of the capacity of the constraining resource, as big enough idle gaps can be filled by internal departures. This results in fewer idle slots and, applied to the constraining resource, results in an increase in throughput. The release of internal departures that match gaps in the arrival stream accurately also results in a more uniform stream, because of the fewer idle gaps. The variability in the meter demand is thus reduced. As discussed in Section 6.3.1 this also results in an increase in the use of available capacity, which as discussed above results in increased throughput. Reduced

variability in demand also results in reduced average delay directly, as discussed in Section 6.3.1.

In addition better scheduling of internal departure release time optimizes the tradeoff between absorbing delay on the ground versus in the air. This results in reduced fuel burn whenever it is beneficial to absorb more delay on the ground.

Modeling of Internal Departure Scheduling

As described in Section 3.2, internal departures enter the system at their actual takeoff time (measured by the first track) rather than at their scheduled departure time. Internal departures were then modeled in a similar fashion to the aircraft arriving en-route. STAs were assigned to the internal departures based on ETAs. The ETAs are calculated based on the average unimpeded transition time to arrive at the first meter fix after departure, and then subsequent meter fixes, as described in Section 3.2. When delay feedback is applied internal departures may absorb a portion of their delay on the ground, delaying their takeoff departure time. Internal departures therefore, were not assessed separately from the rest of the arrival flow and the benefits of including them in the schedule are implicit to the time based metering benefits.

For simplicity and conservatively, modeling internal departures in this benefit assessment does not capture the impact of McTMA on the original takeoff time, which may have already included a delay. As discussed in Section 6.8.4, applying McTMA may have an impact on the original takeoff time because less upstream restrictions may be needed due to the increased capacity by time based metering. Analyzing this impact would require using scheduled rather than actual departure times and estimates of the time spent on the ground, and dealing with the inaccuracies in these estimates, which are still a subject of research for McTMA as explained in Section 5.6. Due to time and resource limitations, it was decided not to analyze the possible change in the original takeoff time and to possibly address it in future research.

7. Estimated McTMA Technical Performance Benefits

In this section the technical benefits of McTMA are outlined along with the details of the associated data analysis. First the data used in the analysis, the assumptions made about key simulation parameters, and the method by which daily benefits were extrapolated to yearly benefits, are described. Then a number of performance metrics are presented, namely: delay savings, throughput increase, fuel burn savings, and metering duration. The daily benefits shown are an average of the results for the fifteen days that were analyzed in the study. These daily benefits were then extrapolated to yearly benefits. Yearly economic benefits were identified accordingly, as detailed in Section 8.

7.1. Data Analyzed

Processed Host and ETMS data were provided by NASA for the facilities of interest (ANY, ZOB, ZBW and ZDC) for a total of 9 days in September 2002 (the 8th, 9th, 10th, 12th, 13th, 14th, 16th, 17th, and 19th) and 30 days in November 2003 (the 1st to 30th). The Host data were in the form of *cm_sim* files, and the ETMS data were in the form of *orig* and *event_list* files. Wind RUC data were also provided by NASA for the same periods. A *Matlab* suite was also provided by NASA for conversion of the *cm_sim* files into a form that could be read by *Matlab*.

Fifteen days were analyzed, all from the November 2003 data set. The days analyzed were November 8th, 9th, 11th, 12th, 14th, 19th, 20th, and 22nd through to 29th. The days were chosen according to data completeness and represented a random and wide range of metering conditions (demand exceeding capacity).

7.2. Simulation Parameters and Assumptions

A key factor in estimating McTMA's benefits is the assumptions made about how McTMA will be used in operation. Certain simulation parameters, representing the modeled benefit mechanisms presented in Section 6, were set according to these assumptions, as detailed below. The assumptions were made according to consultation with NASA's McTMA researchers based on their experience with the tool and in the field.

Duration of Metering

In operation McTMA time based metering is expected to be implemented when demand exceeds runway capacity. The following assumptions are made for deciding when to start and stop metering and which aircraft to meter during these periods:

1. McTMA metering will be implemented when demand exceeds the reported airport capacity, identified in 15 minute windows, as displayed by the Flight Schedule Monitor (FSM). Even though McTMA is not expected to apply the reported capacity of the airport

as a constraint this is the capacity that TMCs are likely to use to identify when demand is high enough to require metering.

2. According to NASA's McTMA researchers McTMA may be turned on in operation about 90 minutes earlier than the time when demand is estimated to exceed the reported capacity at the runway (to capture aircraft as they enter the system and to warm up with a lead period without metering in effect). However, aircraft will not be metered (advised to meet their assigned STAs) until after demand exceeds capacity.

3. As in operation, the start of metering will be identified by a specific aircraft (controllers will decide from which aircraft in the flow to start metering). This aircraft is selected as the first aircraft after demand exceeds capacity for which airport arrival delay is greater than or equal to 2 minutes. The 2 minute threshold was suggested by McTMA researchers and represents an estimate of the amount of delay that can be absorbed in the TRACON. Thus, in the simulation of McTMA delay savings are not counted for aircraft arriving during the lead time until delay under baseline operations reaches 2 minutes.

4. The end time of the McTMA metering period is decided by the end of the queue once no more delay is needed (STA = ETA). However, the end of the metering period should be maintained for sufficient time. Thus, if a new queue starts less than 30 minutes after the end of the previous queue, it is assumed that McTMA is not stopped for such short lull periods in between queues. During these small lull periods delay savings are counted since McTMA is on, due to the reasons explained in 5 below.

5. From the first aircraft in the metering period on all delay savings will be counted until the end of the metering period. Delay savings will be counted even for aircraft delayed due to causes other than the queue (for example, aircraft with STA = ETA because they follow a short lull in the demand within the metering period, but are delayed in the baseline for other reasons). This is because the McTMA researchers believe that the controllers will actively advise all aircraft during the metering period to meet their STAs, in order to avoid inefficiency. Such inefficiency may be caused, for example, by uncertainties in the system (such as internal departures not accounted for in the demand schedule). Not speeding up aircraft behind a lull and leading small queues may contribute to the inefficiency by losing arrival slots that may be used by internal departures not well predicted in advance. Controllers are aware of this and will meter all aircraft in the metering period.[9]

6. Demand is measured over 15 minute periods (because the resolution in the capacity measurement was 15 minutes based on the ASPM database).

7. Departures, in terms of departure rate from the airport, are not taken into account in the decision of when to apply McTMA. Internal departures arriving at the airport are taken into account in the decision of when to apply McTMA, as they add to the arrival demand. Their ETAs are, however, based on their actual departure times, and not scheduled or estimated departure time as may be the case in McTMA operations.

---

[9] Because the baseline model excludes sources of delay other than metering situations (demand exceeding capacity), aircraft that are not delayed by McTMA's time based metering are unlikely to be delayed by the baseline's distance based metering. Therefore, it was observed that the difference in the McTMA delay savings was insignificant whether such aircraft in lull metering periods were assumed to be metered or not.

8. It is assumed that only aircraft metered at the first meter fix are metered at upstream tiers. These aircraft are metered from the time they enter the system. Other aircraft that are within the system when metering starts but not impacted by the queue at the runway are not metered. The assumption is that McTMA will only be used to meter congestion at the runway and will not be used to meter the local congestion in the upstream tiers.
9. McTMA may be applied in operation in conjunction with GDPs. However, because the effect of McTMA operations on GDPs is not analyzed, periods during which GDPs were in place are excluded from the analysis.
10. In McTMA operation flights are unlikely to be metered if the duration of the metering period is very short. The minimum duration of a metering period was chosen as 30 minutes. Metering was thus not applied in the simulation of McTMA for metering period shorter than this.
11. An alternative scenario that was analyzed for the start and stop times of time based metering is that McTMA will be implemented during the same times when MIT was implemented in actual operations. In this scenario the assumption is that McTMA is used merely as a replacement for MIT, which is a conservative assumption. In this scenario the start and end times of McTMA are known. The results for this scenario are presented in Section 9.5 in the sensitivity analysis. In this scenario the times when MIT was applied at the boundary of the system and not propagated from the runway (or TRACON) were excluded. This is because these restrictions are assumed to be due to local constraints while McTMA deals with runway and TRACON constraints.

Capacity parameters - Airport Acceptance Rates (AARs)

McTMA generates STAs based on satisfying an arrival acceptance rate constraint. The arrival capacity was modeled as a function of the departure rate. It was assumed that in each 15 minute period the number of departures is equal to the actual number of aircraft that took off during this period. The arrival rate constraint applied, given this number of departures, was read off the capacity envelopes relating arrival and departure rates described in Section 3.1.1.

It was assumed that McTMA will be limited by runway safety requirements and local operational constraints such as gridlock and controller workload. The capacities imposed on the baseline TRACON model (the capacity envelope percentiles that resulted in the best baseline model calibration as described in Section 4) represent an conservative estimate of this limit, as they represent what capacity is operated under current procedures. McTMA is, however, expected to allow an increase in this capacity due to increased utilization of the available maximum safe throughput limit, as described in Section 6.2. A range of benefits were thus calculated by applying a range of capacities above those that calibrated the baseline model, thus modeling benefits due to an increase in capacity constraints or utilization under McTMA operations. The upper limit of this range of capacities was defined by the maximum safe throughput limit at each airport, and represents what is achievable using McTMA. It does not necessarily represent what controllers will achieve in practice. This will lie somewhere between the current capacity at which the baseline calibrated and the maximum safety limit.

It was assumed that the 99th percentile of the capacity envelopes calculated per half hour (as described in Section 3.1.1) represent the maximum safe throughput limit for each runway configuration. This maximum is determined by the number of runways, the runway configuration and the wake vortex separation requirement. (The 99th percentile is a conservative choice because the 100th percentile may include possible violations of the safety requirements due to controller human error. It is also conservative because it averages the throughput data available in 15 minute intervals, reducing the binning error.) It was assumed that this maximum safety limit will not be increased due to the application of McTMA's time based metering (since McTMA will not impact the runways and wake vortex separations).

Because there is an error in meeting STAs, however, the maximum capacity constraint applied to McTMA must be less than this maximum safe throughput limit. The simulated throughput per 15 minutes, after metering and including the error in meeting STAs, was thus calculated at each airport for a range of capacity constraints (from the percentile of the capacity envelopes that calibrated the baseline TRACON model, to the 98th percentile of the capacity envelopes calculated per quarter hour). Any applied capacity constraint for which the maximum safety constraint (99th percentile of the capacity envelopes calculated per half hour) was not violated by the arrival throughput any more than in actual operations (calculated from actual landing times) was considered to meet the safety requirements. The throughput was permitted to violate the safety constraint to the same extent as actual operations because the 99th percentile of the capacity envelopes calculated per half hour is only an estimate of the safety constraint, and may be conservative in certain cases. It would also be allowing the same level of safety as currently practiced. The highest percentiles of the capacity envelopes that met this criterion are presented in Table 16. The percentage of time (of periods when demand exceeded reported capacity) during which the safety constraints under McTMA and under actual operations were violated are also presented, rounded to the nearest 5 %.

Table 16. Maximum percentile of capacity envelopes calculated per quarter hour that does not violate safety limit any more than under actual operations.

Airport	**$Percentile_{CapEnv}$ (per 15 min) applied to McTMA** [%]	**% time 99th percentile (per 30 min) violated (when Demand > $Capacity_{reported}$)** [%]	
		Actual Throughput	**McTMA Throughput**
PHL	96%	0 %	0 %
LGA	95%	10 %	10 %
EWR	95%	0 %	0 %
JFK	97%	30 %	25 %
TEB	98%	25 %	15 %

It is clear from Table 16 that the safety constraint under actual operations is significantly violated at LGA, JFK and TEB. This suggests that the 99th percentile of the

capacity envelopes calculated per half hour is a conservative estimate of the maximum safe throughput limit at these airports. This is not, however, the case at PHL and EWR. The percentiles of the capacity envelopes applied to McTMA in Table 16 do not allow the throughput under McTMA to violate this limit any more than under actual operations, though, and thus represent a good estimate of the maximum safe capacity limit at each airport, achievable with McTMA.

The range of capacities applied at each airport, as percentiles of the capacity envelopes are presented in Table 17 below. In each case the lower limit represents the capacity at which the baseline TRACON delay model calibrated, and the upper limit represents the maximum safe capacity limit presented in Table 16.

Table 17. Range of capacities applied to McTMA model

Airport	Lower Capacity Limit (Percentile of Capacity Envelope at which Baseline TRACON Delay Model calibrated)	Upper Capacity Limit (Percentile of Capacity Envelope representing maximum safe capacity limit)
PHL	91st	96th
LGA	88th	95th
EWR	88th	95th
JFK	95th	97th
TEB	96th	98th

The upper capacity limit represents what is achievable under McTMA, and is thus used in the nominal case in the results presented in the next sections. Economic results are presented for both limits, however, in Section 8, and benefits for each percentile between these limits is presented in the sensitivity analysis in Section 9.1.

Satisfying Separation Requirements

It was assumed that the acceptance rate constraint is sufficient to allow the satisfaction of the separation requirements at the runway. A 7 Nautical Mile separation requirement was applied at the meter fixes according to consultation with McTMA researchers. This separation was applied to the STAs before applying the error in meeting the STAs.

Internal Departures

Internal departures are modeled in two ways: Departures from major airports are assumed to be assigned a departure time (similar to APREQ or DSP) based on their assigned STA at the next fix downstream. It is assumed that these internal departures takeoff at their actual takeoff times in computing their estimated time of arrival at each airport. Therefore, these internal departures are only delayed relative to their actual time of departure, which might have already been delayed due to an APREQ for example. Applying McTMA may have an impact on the original takeoff time because less upstream restrictions may be needed due to the increased capacity by time based

metering. Analyzing this impact would require using scheduled rather than actual departure times and estimates of the time spent on the ground, and dealing with the inaccuracies in these estimates, which are still a subject of research for McTMA as explained in Section 5.6. Due to time and resource limitations, it was decided not to analyze the possible change in the original takeoff time and to possibly address it in future research.

Departures from smaller airports are not assigned departure times; rather they are captured when they enter the boundary (freeze horizon) of the next tier. An airport was identified as a major internal airport if traffic from this airport constituted more than approximately 1% of arrival traffic into the destination airport under consideration.

Tower En-Route Control (TEC) Traffic

TEC traffic is assumed to arrive at their actual arrival times, reserving time slots. Therefore, their ETAs are assumed equal to their actual times of arrival. They are then assigned STAs and delayed if needed.

Impact on GDPs and Restrictions at System Boundary

Based on consultation with NASA's McTMA researchers it was decided, conservatively and due to time and resource constraints, not to investigate and analyze the impact of McTMA on the other restriction programs, such as GDP, applied in the NAS outside of the McTMA system. This could be done for example by increasing the demand for the McTMA system to take advantage of the increased throughput due to time based metering. Therefore, the arrival and departure demand for the McTMA system was kept unchanged from the current actual level of traffic, as detailed in Section 7.2. The benefits assessment was conducted in periods of time when neither GDP for the analyzed airports nor MIT restrictions on the boundary of the system, were applied.

7.3. Extrapolation to Yearly Benefits

The daily benefits calculated for the 15 days analyzed were extrapolated to yearly benefits for 2003 according to when demand exceeds capacity over the year. Benefits are a function of how often and for how long metering is applied on any given day, and flights are metered under McTMA operations when demand exceeds capacity. Extrapolation of daily benefits from the 15 days analyzed to daily benefits for other days in the year, for which benefits have not been calculated explicitly, can thus be estimated according to how often and for how long demand exceeds capacity. As described in Section 7.2 demand is identified to have exceeded capacity when the demand over any 15 minute period exceeds the reported capacity for the 15 minute period (which is the hourly reported AAR divided by 4). Demand and capacity per 15 minutes is available for every day in the year from the ASPM database. The procedure by which the daily results for the 15 days analyzed is extrapolated to a year is thus as follows:

1. The percentage of the day for which demand exceeded reported capacity was identified for each day in 2003, according to data from the ASPM database.
2. Every day in 2003 was then categorized according to this percentage. Four categories were identified for the analysis. These categories are evenly distributed over the range of percentages of the day for which demand exceeded capacity. Only 4 categories were used to ensure that at least one analyzed day fell into each category, from which the average daily delays could be calculated for each category.
3. The number of days in the year falling into each category was calculated.
4. The 15 days analyzed were categorized in the same categories.
5. The average daily benefits for each category were then calculated according to the daily benefits of the analyzed days that fell into each category.
6. Daily benefits were then extrapolated to yearly benefits by multiplying the average daily benefits for each category by the number of days in the year in each category.

In the sensitivity analysis case for which aircraft are only metered by time based metering when MIT were applied, daily benefits are extrapolated to yearly benefits according to when MIT restrictions were implemented at the TRACON boundary, and not when demand exceeded capacity. This is because this is more representative of the cause of metering in this scenario. Because data for when MIT were applied is only available for the month of November 2003, the extrapolation is based on November only. The yearly results are then calculated assuming the same percentage of MIT impacted periods for the rest of the year as identified for November. The results for this analysis are presented in the sensitivity analysis in Section 9.5.

7.4. Delay Savings

Baseline delay was calculated for each arrival flight at each of the five airports under consideration, for each of the fifteen days analyzed. Each flight's baseline arrival delay was calculated as the flight's modeled time of arrival at the airport minus its ETA. Average arrival delay per flight during baseline metering periods over the fifteen days analyzed are presented in Table 18 below for each airport, alongside the average number of flights impacted by baseline metering per day. These results are the average of the results for each day individually.

Arrival delay was also calculated for each arrival flight in the simulation of McTMA at each of the five airports under consideration, for each of the fifteen days analyzed. Each flight's McTMA arrival delay was calculated as the flight's modeled time of arrival (STA plus error), minus the ETA. Average arrival delays per flight during McTMA metering periods over the fifteen days are also presented in Table 18 for each airport, alongside the average number of flights impacted by McTMA metering per day. The results apply the percentile of the capacity envelopes representing maximum safe capacity limit (increased capacity). This is the expected achievable increase in capacity over baseline operations as described in Section 7.2.

It is important to note that because baseline and McTMA operations are metered during different periods, the average delays per flight presented in Table 18 are not averages over the same flights.

Table 18. Baseline and McTMA delay per flight.

Airport	Baseline		McTMA Increased Capacity	
	Average Delay [min / flt]	Average Number of Flights Metered [flts / day]	Average Delay [min / flt]	Average Number of Flights Metered [flts / day]
PHL	10.293	195	5.367	145
LGA	8.072	190	4.192	193
EWR	8.394	124	2.773	107
JFK	10.276	41	4.922	41
TEB	4.834	21	1.976	15

In Table 18 average delay during metering under McTMA is less than that under baseline operations, at each of the airports. This means that during periods of metering, McTMA is able to meet the applied AAR while incurring less delay per flight than in baseline operations. This reduction in delay per flight only accounts for reductions in delay due to the use of time based metering instead of MIT and TRACON metering, and not due to weather.

JFK and PHL show the highest average delay per flight under McTMA metering of all the airports studied. This is likely to be because they are the only airport studied where significant banks are scheduled. During banks, demand generally exceeds capacity by a significant amount, and aircraft incur significant delay even under McTMA metering. With a flatter schedule, demand does not exceed capacity to the same extent, and thus the delay during metering is less. Note that LGA, which operates a flatter schedule, but very close to capacity for much of the day, shows a higher average number of flights affected by metering per day than PHL, but has a lower average delay per day. EWR and TEB have lower average delay per day, but also fewer flights affected by metering per day.

Figure 90 shows an example of the reduction in delay between McTMA and baseline operations.

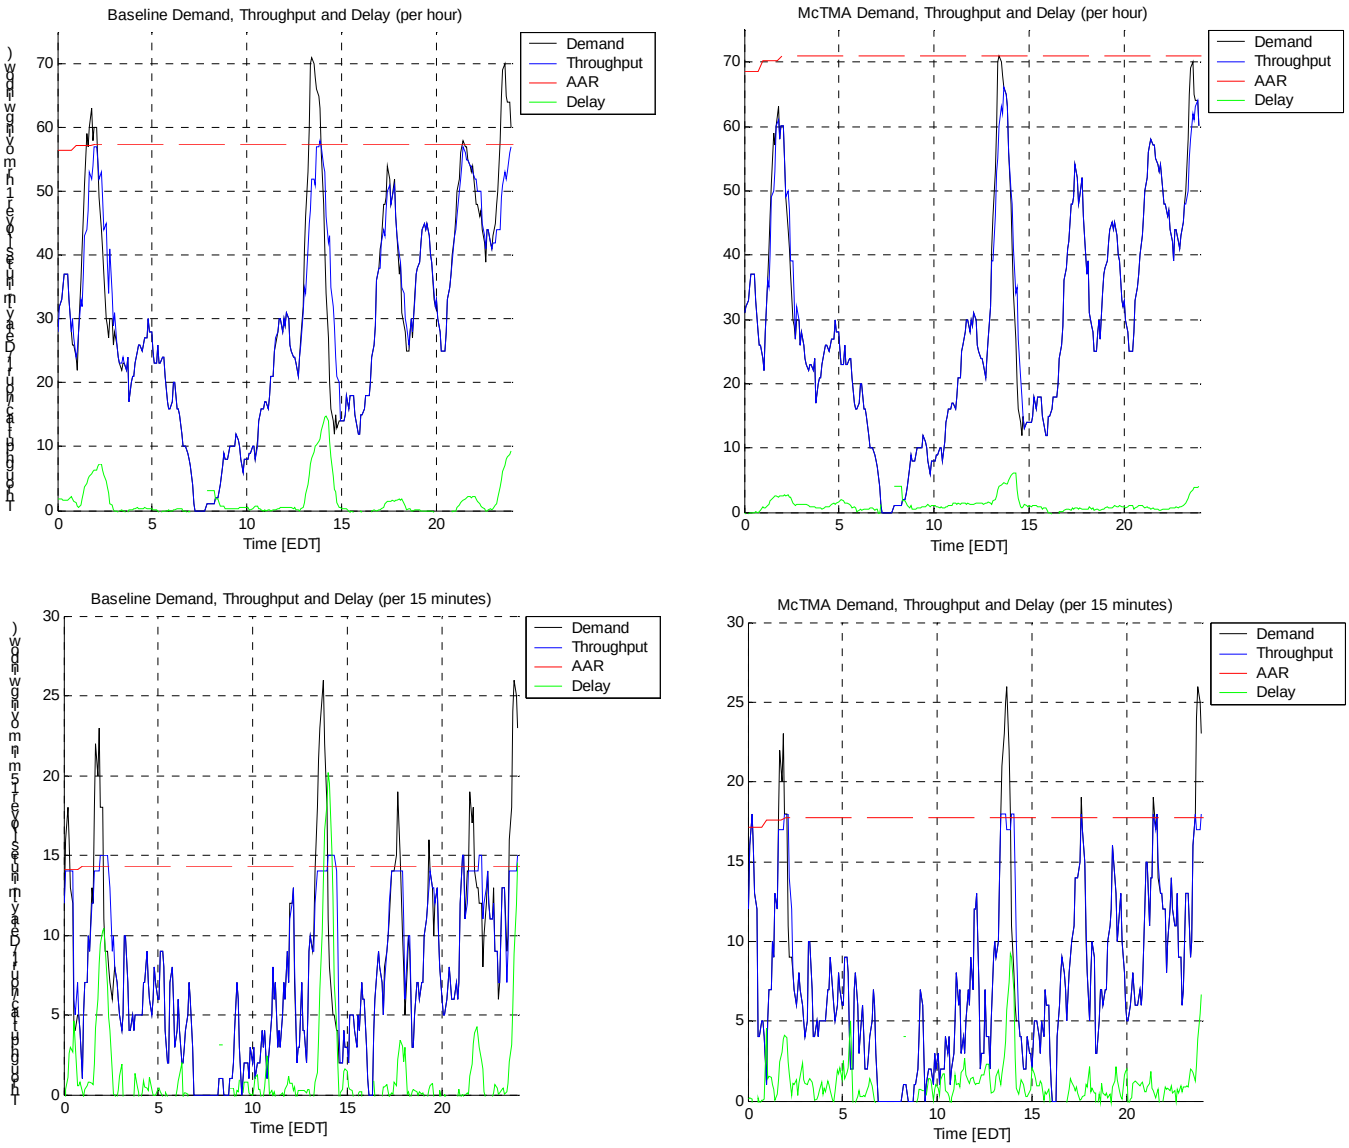

Figure 90. Throughput and delay plotted for PHL for Nov. 25, 2003, under current operations, and under McTMA. The upper two plots show results per hour, and the lower two show results per 15 minutes.

It is clear from Figure 90 that the delay under McTMA operations (the right two plots) is lower than under baseline operations. These plots also show how each model restricts the throughput to the applied capacity. In the McTMA cases the 96th percentile of the capacity envelope is applied, while in the baseline case the 91st percentile of the capacity envelope is applied.

Delay Savings per Year

The total delay savings, extrapolated to the year (2003) as described in Section 7.3, are presented in Table 19 below. The applied capacity is that assuming the expected achievable increase in capacity as described in Section 7.2.

Table 19. McTMA delay savings per year (2003) assuming an increase in capacity under McTMA.

Airport	Yearly McTMA Delay Savings over Baseline Operations **2003** [min]
PHL	395,000
LGA	267,000
EWR	263,000
JFK	51,000
TEB	10,000

The highest delay savings result at PHL, LGA, and EWR, with JFK and TEB significantly lower. This is because JFK and TEB have lower traffic levels, and are less constrained under current operations, as described in Section 3.1.1. The introduction of McTMA to these airports is not thus likely to have as significant an impact on delays, unless the traffic levels are increased closer to saturation.

With no increase in capacity the total delay savings are as detailed in below.

Table 20. McTMA delay savings per year (2003) assuming no increase in capacity.

Airport	Yearly McTMA Delay Savings over Baseline Operations **2003** [min]
PHL	24,000
LGA	20,000
EWR	22,000
JFK	17,000
TEB	0

The delay savings with no increase in capacity are clearly significantly lower than those with increased capacity. This is because, with no increase in capacity, the delay savings with McTMA are exclusively from improved metering efficiency. This suggests that significant benefits are gained by allowing the airport capacity to increase to utilize the reduced system service rate variability under McTMA.

TEB particularly shows no significant delay savings. This is likely to be because of the very low demand at this airport, requiring very little metering.

Distribution of Delay savings over Tiers

The average delay absorbed in each tier under modeled baseline and McTMA operations were calculated over the 15 days analyzed, and compared. The results for PHL are presented in graphical form in Figure 91 below. The results for the other airports show similar trends.

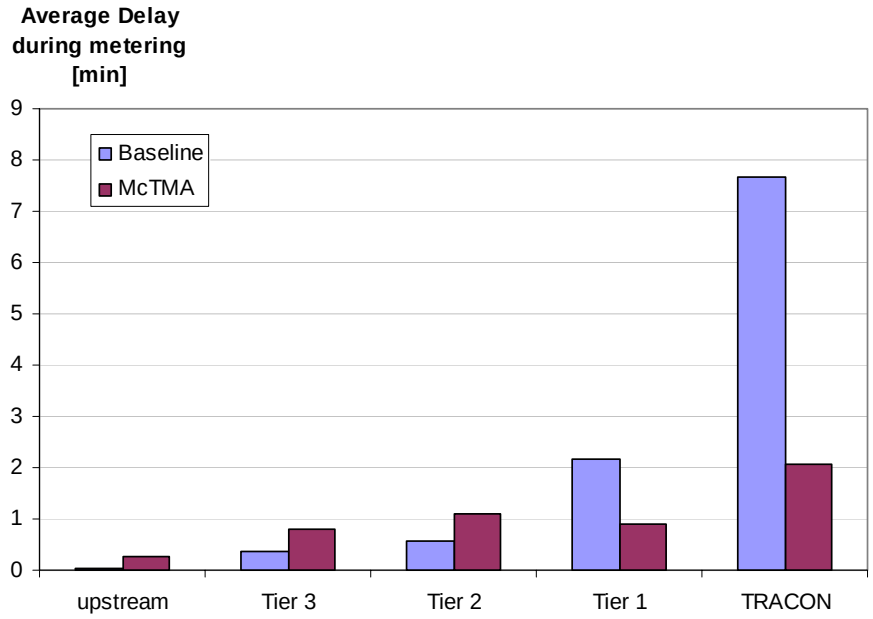

Figure 91. Average delay per tier under modeled baseline and McTMA operations, EWR

It can be seen from the results presented in Figure 91 that under modeled McTMA operations, the delay is distributed more evenly over the four tiers than under modeled baseline operations. This is because the delayabilities within the McTMA model force delay to be fed upstream when no more delay can be absorbed between any meter fix pair without holding. The distribution of delay upstream can be adjusted by adjusting the delayabilities in each tier. If the delayabilities in the TRACON were reduced, thus requiring the TRACON to absorb less of the delay that must be absorbed, the average TRACON delay under McTMA would be reduced in Figure 91, and the average delay in the upstream tiers increased. In this way delay can be distributed to the ideal tier for reduction in fuel burn. In this analysis delayability has not been adjusted specifically to minimize fuel burn. Instead the delayability is simply the maximum delay that can be absorbed without holding, as described in Section 3.1.3. The sensitivity of the economic benefits to delayability is presented in Section 9.2.

7.5. Throughput Increase

The increase in throughput resulting from McTMA operations is related to both improved arrival flow efficiency, and the increased airport capacity applied under modeled McTMA operations. The measurement of throughput is, however, highly sensitive to the period over which the throughput is measured. The results calculating the increase in throughput for those periods when flights were metered under McTMA only are presented in Table 21 below. The applied capacity is that assuming the expected achievable increase in capacity as described in Section 7.2.

Table 21. Increases in throughput during McTMA metering (2003) assuming an increase in capacity.

Airport	Increase in Throughput during McTMA metering periods 2003 [%]
PHL	4.96 %
LGA	2.10 %
EWR	2.61 %
JFK	2.87 %
TEB	0.00 %

The increase in throughput is highest at PHL, and lowest at TEB (negligible). The throughput increase is likely to be highest at airports that operate significant arrival banks, because the demand must be particularly high to utilize an increase in capacity over an extended length of time. PHL operates such a significant bank structure. At airports where demand does not exceed capacity to the same extent, such as EWR, JFK and LGA, the arrival queues are shorter. Calculating throughput over metering periods which include a number of queues thus averages out much of the increase in throughput in each queue. TEB likely shows a negligible increase in throughput because demand does not often extend above capacity.

With no increase in capacity the increases in throughput due to McTMA at each of the airports are presented in Table 22 below.

Table 22. Increases in throughput during McTMA metering (2003) assuming no increase in capacity.

Airport	Increase in Throughput during McTMA metering periods 2003 [%]
PHL	2.41 %
LGA	1.31 %
EWR	2.25 %
JFK	2.56 %
TEB	0.00 %

With no increase in capacity, the increase in throughput due to McTMA operations is lower at each of the airports. The change is greatest at PHL. This is likely to be because of the banking structure operated, which means the marginal benefit of any increase in capacity is higher. This effect is discussed in detail in Section 9.1. LGA, EWR, JFK and TEB show little change in throughput increase. This is likely to be because these aiports do not operate significant banking structures, like at PHL.

7.6. Fuel Burn Savings

Savings in fuel burn were also calculated for each arrival flight at each of the five airports under consideration, for each of the fifteen days analyzed. Fuel burn savings were estimated by calculating the amount of delay absorbed in each tier of the arrival network, for each flight, for both the modeled baseline case, and the modeled McTMA case. The delay in each tier was then multiplied by an average fuel burn rate (lb per hour) according to the aircraft's speed, altitude, and weight class, to yield the total fuel burned in delay, for each flight. Speed and altitude for the upstream fix in each tier were used for the analysis because delay is generally absorbed during the cruise segment of the flight, before the decent segment, which would occur in the end portion of the tier. Average fuel burn rate according to aircraft speed, altitude, and weight class was obtained from NASA CTAS engineering [9], assuming a cruise flight phase. This is consistent with the assumption that delay is generally absorbed during the cruise segment of the tier, and not during the decent segment. Figure 92 shows the Mach/altitude profile identified for Large aircraft types, and the corresponding fuel burn/altitude profile obtained from NASA CTAS engineering. The Mach/altitude profile is an average calculated from an analysis of host track data for September 10, 17 and 19, 2002. Mach/altitude profiles were also identified for Small, Heavy and B757 aircraft types, and corresponding fuel burn/altitude profiles obtained from NASA CTAS engineering.

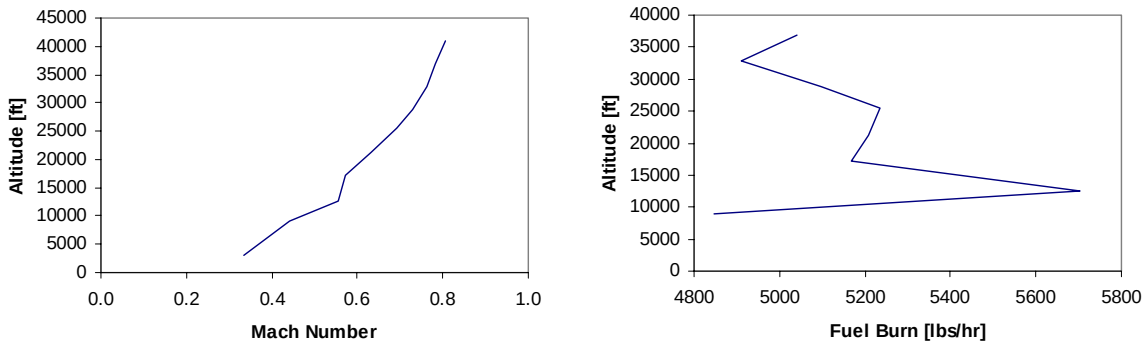

Figure 92. Mach/altitude profile and corresponding fuel burn/altitude profile for Large aircraft types.

It is clear from the fuel burn/altitude profile that fuel burn does not decrease monotonically with altitude. This is unexpected, and may be because Mach/altitude profiles vary significantly by type, even within the Large aircraft type class. The other aircraft types show similar effects.

In order to calculate fuel burn savings, fuel burned during delay under McTMA was subtracted from the fuel burned during delay in the baseline case, for each flight, yielding a fuel burn saving (in lb) for each flight. These results were then extrapolated to yearly savings as described in Section 7.3, and are presented in Table 23 below. The applied capacity is that assuming the expected achievable increase in capacity as described in Section 7.2. As a reference for the order of magnitude of these savings, an *Airbus A320* burns as much as 32,000 lbs of fuel on a transcontinental flight.

Table 23. Average fuel burn savings per year (2003) assuming an increase in capacity under McTMA.

Airport	**Average Fuel Burn Savings 2003** [lb]
PHL	19,005,000
LGA	15,266,000
EWR	15,873,000
JFK	5,741,000
TEB	520,000

Relative to the amount of fuel burned on a single flight, the fuel burn savings resulting from McTMA per year are high (in the order of 650 transcontinental flights, at PHL). These savings are significant as fuel costs are one of the primary components of airline direct operating costs.

It can be seen that the fuel burn savings are highest at EWR, LGA and PHL, and lowest at JFK and TEB. This is because JFK and TEB have lower traffic levels, and are less constrained under current operations, as described in Section 3.1.1.

With no increase in capacity the fuel burn savings are as detailed in Table 25 below.

Table 24. Average fuel burn savings per year (2003) assuming no increase in capacity.

Airport	Average Fuel Burn Savings 2003 [lb]
PHL	2,850,000
LGA	2,613,000
EWR	3,066,000
JFK	3,041,000
TEB	33,000

The fuel burn savings with no increase in capacity are clearly significantly lower than those with increased capacity, as expected.

7.7. Duration of Metering

The duration of each metering period under modeled McTMA metering was calculated for the 15 days analyzed, and the minimum, maximum and mean durations identified. As described in Section 7.2 the minimum duration for metering was restricted to 30 minutes. Any metering duration less than this threshold was not counted. The mean, maximum and minimum durations of McTMA metering periods are presented in Table 25. The applied capacity is that assuming the expected achievable increase in capacity as described in Section 7.2.

Table 25. Duration of McTMA metering.

Airport	Mean McTMA Metering Period Duration [min]	Maximum McTMA Metering Period Duration [min]	Minimum McTMA Metering Period Duration [min]
PHL	52	345	30
LGA	109	600	30
EWR	66	345	30
JFK	59	165	30
TEB	42	105	30

According to the results presented in Table 25 the mean McTMA metering period duration is significantly higher than the 30-minute applied minimum duration at each airport. LGA shows the highest average of 1 hr 49 minutes, while TEB shows the lowest of 42 minutes. This is because LGA operates particularly close to capacity for much of the day, and is thus expected to require longer periods of metering. TEB, however, does not operate close to capacity as often, and does not operate clear banks, so is not expected to require long periods of metering. Although JFK has lower traffic, the mean McTMA metering period duration is high, at 59 minutes. This is likely to be because although JFK does not violate capacity often, it does operate significant banks (particularly in the afternoon), which are likely to require significant metering. Thus, when the airport requires metering, the duration of the metering is long.

LGA shows the highest maximum duration of metering, of 10 hrs. LGA operates at high demand throughout the day with very few lulls in demand, and thus has little opportunity for delay recovery under low capacity. When capacity is low metering is thus likely to be required through most of the day. JFK and TEB show maximum durations of only 2 hrs 45 minutes and 1 hr 45 minutes respectively. This is because of the lower demand at these airports.

The minimum McTMA metering period duration was 30 minutes in all cases, which corresponds to the applied minimum duration. This means the minimum duration was a constraining factor at all airports.

8. Estimated McTMA Economic Benefits

The benefits of McTMA identified from the analysis of the benefit mechanisms in Section 6 are increased throughput, reduced average delay and reduced fuel burn as presented in Section 7. The mechanisms by which these benefits result in direct economic benefits are detailed in Figure 93.

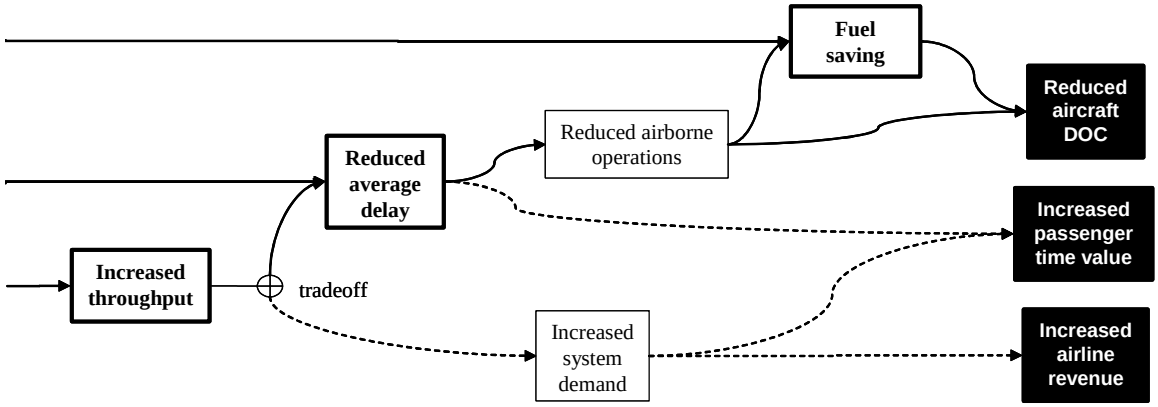

Figure 93. Direct economic benefits.

Without any strategic responses by the airlines to change demand, increased throughput results directly in reduced average delay. A reduction in average delay results directly in a reduction in airborne operations, and thus aircraft direct operating costs. However, increased throughput also allows the airlines to increase demand. An increase in demand increases revenue according to the operating profit made on each of the flights added, but also increases average delay. There is thus a tradeoff between increasing demand and reducing delay. An increase in demand requires a strategic response by the airlines and is not modeled in this study. This study thus considers the benefits associated with a reduction in delay only.

Reduced delay and increased demand also result in added value to passengers in terms of time savings and increased travel time opportunities. These however are also not modeled in this study, based on feedback from NASA's researchers.

8.1. Methodology for Calculation of Economic Benefits

Economic Benefits of Delay Savings

Average aircraft operating costs per airborne hour were estimated according to the FAA APO 1998 [10], for each of the 4 different aircraft weight classes (small, large, heavy and Boeing B757). The results were corrected for inflation from 1996 Dollars to 2003 Dollars according to an inflation calculator provided by the U.S. Department of Labor Bureau of Labor Statistics [11]. Because fuel burn savings are calculated separately in the simulations, fuel and oil costs were excluded from the values estimated. The cost of the delay (excluding fuel) of each flight was thus calculated according to the minutes of delay relative to the flight's ETA, and according to the aircraft's weight class.

Economic Benefits of Fuel Burn Savings

Fuel cost was estimated for November 2003, as recorded by the Bureau of Transportation Statistics Office of Airline Information [12], at US$ 0.84 per gallon. The economic savings resulting from the savings in fuel burn were calculated using this value.

Daily Economic Benefits

The economic benefits for each day were calculated by comparing the total operating profit in the baseline case to the total operating profit with McTMA in operation. The baseline operating profit per flight can be estimated according to historic financial data recorded by the Bureau of Transportation Studies Office of Airline Information [13]. As presented in equation (11) below, dividing system operating profit by total system departures, yields operating profit per flight.

$$Operating\ Profit\ per\ flt\ [US\$/flt] = \frac{Operating\ Profit\ [US\$/quarter]}{Total\ Operations\ [Deps/quarter]} \qquad (11)$$

Operating profit is calculated by subtracting operating cost from operating revenue. Operating profit per flight can thus be calculated by subtracting operating cost per flight from operating revenue per flight. In order to exclude cyclical effects, the average operating cost and revenue per flight, for all quarters from the first quarter of 1996 to the third quarter of 2003 were calculated. Average operating cost was found to be US$13,546 per flight (with a standard deviation of US$ 1,108 per flight). Average operating revenue was found to be US$ 13,888 per flight (with a standard deviation of US$ 1,660 per flight). This yields an operating profit of US$ 342 per flight. Baseline operating profit on each day can be estimated accordingly by multiplying by the number of flights on each day.

The operating profit with McTMA in operation can be estimated by adding the economic savings due to delay and fuel burn, per flight, to the baseline operating profit of US$ 342 per flight. Operating profit per day, using McTMA, can thus be estimated according to equation (12) below.

$$\begin{aligned} McTMA\ Operating\ Profit\ per\ day\ [US\$] \;=\; & Total\ Operations\ per\ day\ [Arrs] \times \\ & (Operating\ Profit\ per\ flt\ [US\$/flt] + Delay\ Savings\ per\ flt\ [US\$/flt] + \\ & \qquad Fuel\ Burn\ Savings\ per\ flt\ [US\$/flt]) \end{aligned} \qquad (12)$$

Total increase in operating profit for the day, due to McTMA, can then be calculated by finding the difference between the baseline and McTMA operating profits for the day. In this calculation operating profit per flight cancels out.

Extrapolation to Yearly Benefits

The daily economic benefits are extended to results for the entire year (2003) according to the process described in Section 7.3, extrapolating benefits according to when demand exceeds capacity.

8.2. Yearly Economic Benefit Results

The results of the economic benefit analysis described in Section 8.1, which describe the yearly increase in operating profit due to McTMA operations over current baseline operations, are presented in Table 26 below. The applied capacity is that assuming the expected achievable increase in capacity as described in Section 7.2.

Table 26. Yearly increase in operating profit due to McTMA (2003) assuming an increase in capacity under McTMA.

Airport	2003 Yearly Savings [US$ / year]
PHL	17,275,000
LGA	11,810,000
EWR	12,817,000
JFK	3,289,000
TEB	383,000

The highest benefits per year with an increase in capacity are seen at PHL, EWR, and LGA. As with delay savings (Section 7.4) and fuel burn savings (Section 7.6) JFK and TEB show lower benefits. This is because there is less traffic at these airports and the demand does not saturate the airport capacity to the degree that it does at LGA, EWR and PHL. It is thus expected that McTMA will show lower benefits at these airports.

The benefits at PHL are quite significantly higher than those at EWR and LGA. This is because the percentile of the capacity envelope applied at PHL, representing increased capacity (96^{th} percentile), is higher than those applied at EWR and LGA (95^{th} percentile); and because the marginal benefit of increased capacity is higher at PHL than EWR and LGA. The marginal benefit of capacity at each airport, which is related to the bank structure operated at each airport, is discussed in detail in Section 9.1 in the discussion on the sensivity of the result to capacity.

The benefits of McTMA assuming no increase in capacity are presented in Table 27 below.

Table 27. Yearly increase in operating profit due to McTMA (2003) assuming no increase in capacity.

Airport	**2003 Yearly Savings** [US$ / year]
PHL	1,302,000
LGA	1,141,000
EWR	1,343,000
JFK	1,268,000
TEB	4,000

With no increase in capacity assumed, which represents the lowest limit for McTMA benefits, all airports show significantly lower benefits than with an increase in capacity. This suggests that significant benefits are gained by allowing the airport capacity to increase to utilize the reduced system service rate variability under McTMA. However, without this increase McTMA does still show benefits (except at TEB). These benefits are exclusively due to improved metering efficiency and feeding delays upstream to higher altitudes where the fuel burn rate is reduced.

9. Sensitivity Analysis

There are two reasons for which parameters were varied and the model sensitivity tested. The first reason is to vary parameters that represent benefit mechanisms, to push the limits of the system under McTMA. The second is to vary parameters to test whether or not the developed models hold under varying assumptions about the values of these parameters. The following parameters were varied to test benefit mechanisms:

- Airport capacity under McTMA, represented by the airport acceptance rates (AARs) from the airport capacity envelopes.
- Sector capacities, represented by each sector's delayability, or amount of delay that can be absorbed in each sector without holding.
- Error ε in meeting STAs.
- Operating cost per flight.
- Metering periods, including the minimum duration of metering, and the criteria for initiating metering.

The results of varying each of the above parameters is discussed in detail below. The nominal case presented in each section applies the expected achievable increase in capacity under McTMA operations, as described in Section 7.2.

9.1. Airport Capacity under McTMA

The airport capacity applied under McTMA was varied over a range of values to investigate the sensitivity of the benefits to this parameter. The lowest capacity applied was the percentile of the capacity envelope that was modeled to constrain the baseline TRACON model, while the highest capacity applied was the 98th percentile of the capacity envelope. Applying the percentile of the capacity envelope that was modeled to constrain the baseline TRACON model constrains McTMA to the same airport capacity as the baseline, thus modeling the benefits of McTMA if no additional airport capacity over current operations is utilized. This assumes that under McTMA each airport's capacity is underutilized to the same extent as it is under current baseline operations. This is a particularly conservative assumption because McTMA allows an increase in applied capacity because of the reduced variability in the system service rate, as explained in Section 6. The identification of the percentile of the capacity envelope that constrains the baseline TRACON model is described in Section 4.2, under calibration of the baseline TRACON delay model. The percentiles of the capacity envelopes that constrain the baseline TRACON models were found to be as follows: PHL - 91st percentile, LGA – 88th percentile, EWR – 88th percentile, JFK – 95th percentile, TEB – 96th percentile.

The results presented in Table 28 are significantly lower than presented in Section 8.2. This suggests that significant benefits are gained by allowing the airport capacity to increase under McTMA. However, without this increase McTMA does still show benefits. These benefits are due to improved metering efficiency and feeding delays upstream to higher altitudes where the fuel burn rate is reduced.

As described in Section 3.1.1 the 99th percentile of the capacity envelope calculated per half hour was considered to represent the maximum safe limit to capacity. The 98th percentile of the capacity envelope was thus chosen as the highest capacity applied in the sensitivity analysis. Table 28 below shows the yearly benefits for each of the AARs modeled, from the percentile of the capacity envelope that was modeled to constrain the baseline TRACON model, to the 98th percentile of the capacity envelope. This provides a range of the benefits resulting from McTMA depending on what airport capacity is applied. These results are presented in graphic form in Figure 94 below. The nominal case, which represents the capacity increase achievable without violating the maximum safe capacity limit, is highlighted in bold in Table 28, and also marked on the curves in Figure 94.

Table 28. Yearly increase in operating profit due to McTMA, 2003 – varying applied airport capacity under McTMA.

Percentile of Capacity Envelope	**2003 Yearly Savings** [US$ / year]				
	PHL	**LGA**	**EWR**	**JFK**	**TEB**
88th %		1,141,000	1,343,000		
89th %		2,478,000	3,602, 000		
90th %		3,353,000	5,505,000		
91st %	1,302,000	5,623,000	7,699,000		
92nd %	5,809,000	5,946,000	9,095,000		
93rd %	9,855,000	8,838,000	10,547,000		
94th %	13,112,000	11,045,000	11,938,000		
95th %	15,677,000	**11,810,000**	**12,817,000**	1,268,000	
96th %	**17,275,000**	12,691,000	13,430,000	2,579,000	4,000
97th %	19,219,000	14,071,000	13,358,000	**3,289,000**	157,000
98th %	19,970,000	15,538,000	13,819,000	4,435,000	**383,000**

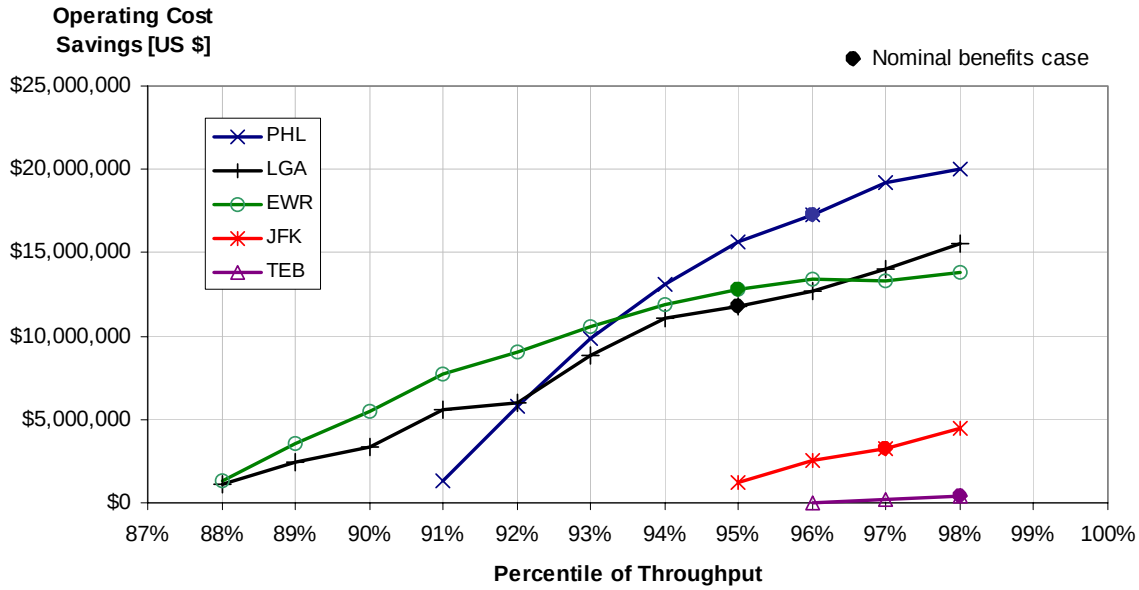

Figure 94. Yearly increase in operating profit due to McTMA, 2003 – varying applied airport capacity under McTMA.

Its is clear from the results in Table 28 and from Figure 94 that the economic benefits of McTMA increase with increasing airport capacity. The rate of increase, however, consistently decreases at higher percentiles of the capacity envelopes. This is because as the airport capacity increases, the relative benefit of the increase decreases, as there are fewer periods of demand that are able to utilize the increased capacity. The benefits are expected to reach a maximum above which any further increase in capacity would produce no further benefit, because all the demand is under capacity.

The curve for PHL is steeper than those for LGA and EWR. This means that the marginal benefits from increasing the capacity by one percentile are larger at PHL than the other airports. The larger steepness at PHL is likely to be due to the schedule bank structure operated at PHL, and the higher demand to capacity ratio during the banks. At LGA and EWR the demand is more spread out and thus, after few percentile increases in capacity, the benefits plateau. This is because the capacity satisfies the demand at these percentiles. At PHL it takes more percentile increases in capacity to exceed the demand because of the large demand peaks. The marginal benefits for each percentile increase in capacity are thus higher.

It is also clear from Figure 94 that the increase in benefits is smooth in all cases except LGA. The lack of smoothness at LGA is due to resolution problems in the identification of the 91st to 96th percentiles of the capacity envelopes for LGA. As can be seen in Figure 9 there are portions of the capacity envelopes at LGA for which resolution between different percentiles is limited, as discussed in detail section 3.1.1.

9.2. Delayability

The delayability between each meter fix pair was also varied, and the yearly benefits recorded. Delayability was calculated as the 90th percentile of delay absorbed between meter fixes, relative to the unimpeded transition times between meter fix pairs, for the results presented in Sections 7 and 8. The resulting delayability was increased and decreased by 30%, and the benefits calculated, to investigate the sensitivity to this parameter. The results are presented in Table 29 below. Such an increase (30%) in delayability corresponds approximately to using the 95th percentile of delay absorbed between meter fixes for calculation of the delayability, instead of the 90th percentile.

Table 29. Yearly increase in operating profit due to McTMA, 2003 – varying Delayability.

Airport	2003 Yearly Savings [US$ / year]		
	Nominal Case	Increase in Delayability by 30%	Decrease in Delayability by 30%
PHL	17,275,000	17,481,000	17,628,000
LGA	11,810,000	11,959,000	12,069,000
EWR	12,817,000	12,957,000	12,857,000
JFK	3,289,000	3,391,000	3,369,000
TEB	383,000	352,000	358,000

It is clear from the results presented above that the changes in yearly savings resulting from increases and decreases in delayability are small. The changes in the results are small enough to be in the noise caused by the variabilities modeled in the baseline and McTMA models, so a trend cannot be identified. It is expected that a decrease in delayability would increase the benefits as lower delayabilities, particularly at the downstream tiers, increase the amount of delay fed back upstream. If fuel burn is lower at higher altitudes, economic savings would be expected to increase with this effect. As discussed in Section 7.6, however, the relationship between fuel burn and altitude is not simple.

9.3. Error in Meeting Scheduled Times of Arrival

The error is meeting STAs was modeled by sampling from a normal distribution centered at zero, with a standard deviation of 90 sec, and a maximum and minimum error of two standard deviations (180 sec and -180 sec respectively). According to a NASA study at DFW, however, the standard deviation in the error in meeting STAs is 150 sec. The mean error in this study was found to be very close to zero. The standard deviation was thus increased to 150 sec, and decreased to 60 seconds to determine the sensitivity of

the results to the variability with which aircraft meet their STAs. The maximum and minimum errors were also adjusted accordingly (two standard deviations). The effect of increasing the mean error from zero to 60 seconds was also analyzed, modeling a bias towards flights arriving at the meter fixes late, on average. Similarly the effect of decreasing the mean error from zero to -60 seconds was also analyzed, modeling a bias towards flights arriving at the meter fixes early, on average. According to McTMA researchers aircraft are likely to arrive at meter fixes early more often than late because controllers are in compliance even if aircraft arrive at their meter fixes when the delay request displayed to them is 1 minute, and not only when it is zero. This means that an aircraft may be 1 minute early and still be in compliance.

The results for each of these cases is presented in Table 30 below.

Table 30. Yearly increase in operating profit due to McTMA, 2003 – varying the mean error in meeting STAs (μ) and standard deviation in meeting STAs (σ).

Airport	2003 Yearly Savings [US$ / year]				
	μ = 0 sec σ = 90 sec (Nominal)	**μ = 0 sec σ = 150 sec**	**μ = 0 sec σ = 60 sec**	**μ = 60 sec σ = 90 sec**	**μ = –60 sec σ = 90 sec**
PHL	17,275,000	16,527,000	17,782,000	13,031,000	21,357,000
LGA	11,810,000	10,435,000	12,538,000	4,842,000	17,917,000
EWR	12,817,000	11,595,000	13,549,000	8,721,000	16,414,000
JFK	3,289,000	3,275,000	3,551,000	1,950,000	4,762,000
TEB	383,000	280,000	380,000	135,000	516,000

The effect of varying the standard deviation in the error in meeting STAs can be seen in Figure 26 to be small. As the standard deviation in the error in meeting the STA is increased savings decrease slightly, and as it is reduced, savings increase slightly. This is expected as increased variability in meeting STAs is expected to reduce the benefits of McTMA. It is clear, however, that the benefits are not sensitive to the standard deviation in the error in meeting STAs.

The effect of increasing the mean error in meeting STAs to 60 seconds can be seen in Figure 26 to be significant, reducing the benefits of McTMA quite severely. This is because the increased mean error in meeting STAs forces aircraft to be late on average, which increases delays under McTMA operations. This suggests that it is particularly important to ensure that the mean error in meeting the STAs is not greater than zero.

The effect of decreasing the mean error in meeting the STAs to -60 seconds can be seen in Figure 26 to increase savings significantly. This is expected as this analysis

models aircraft to arrive early, on average, which reduces average delay. This increase in savings assumes that all errors in meeting STAs are centered at -60 seconds. This includes flights for which STAs are equal to ETAs. For these flights, this would require the aircraft to arrive at the meter fixes earlier even than its ETAs. This is not unreasonable, however, as the average standard deviation in unimpeded transition times for each airport, which define flight ETAs, is 90 seconds. This means that an aircraft arriving 90 seconds early has transition times within approximately one standard deviation of estimated unimpeded transition times.

9.4. Operating Cost per Flight

As detailed in Section 8 above, operating cost per flight was estimated as the average operating cost per flight for all quarters from the first quarter of 1996 to the second quarter of 2003. The standard deviation in operating cost over this period was US$ 1,108 per flight, which is 8% of the average. Operating cost per flight, for both McTMA and baseline operations, was thus increased and decreased by this standard deviation, and the yearly benefits recalculated, as presented in Table 31. This results in a corresponding 8% increase and decrease in benefits.

Table 31. Yearly increase in operating profit due to McTMA, 2003 – varying operating cost per flight.

Airport	2003 Yearly Savings [US$ / year]		
	Nominal Case	8% increase in Operating Cost per flight	8% decrease in Operating Cost per flight
PHL	17,275,000	18,657,000	15,893,000
LGA	11,810,000	12,755,000	10,865,000
EWR	12,817,000	13,842,000	11,792,000
JFK	3,289,000	3,552,000	3,026,000
TEB	383,000	414,000	352,000

The results in Table 31 suggest that McTMA will show high benefits regardless of the state of the airline industry. This is because McTMA reduces costs significantly, which otherwise remain relatively constant. It does not reduce revenue, which is highly dependent on the economy and state of the airline industry.

9.5. Metering Periods

The sensitivity of the economic benefits of McTMA to the metering periods, including the sensitivity to the minimum duration of the metering periods, and the criteria whereby metering is initiated, are presented below.

Minimum Duration of Metering

As discussed in Section 7.2, the minimum duration for metering was chosen as 30 minutes. Alternative values of 15 minutes and of 60 minutes were analyzed, and the sensitivity of the benefits to this parameter identified. The results are presented in Table 32 below.

Table 32. Yearly increase in operating profit due to McTMA, 2003 – varying minimum duration of metering.

Airport	2003 Yearly Savings [US$ / year]		
	Nominal Case Minimum Duration = 30 min	Minimum Duration = 15 min	Minimum Duration = 60 min
PHL	17,275,000	18,313,000	11,970,000
LGA	11,810,000	11,590,000	10,024,000
EWR	12,817,000	14,170,000	10,156,000
JFK	3,289,000	3,574,000	2,977,000
TEB	383,000	397,000	268,000

According to the results presented in Table 32, the sensitivity to the minimum duration of metering is small in most cases. Benefits are generally increased when the minimum duration of metering is reduced to 15 minutes, because more periods are metered (fewer periods are discarded because they are too short). The changes in the results are small enough to be in the noise caused by the variabilities modeled in the baseline and McTMA models. The benefits are consistently lower when the minimum duration of metering is increased to 60 minutes. This is because fewer periods are metered (more periods are discarded because they are too short), and thus the benefits of metering are lower.

Criteria for Initiating Metering

As discussed in Section 7.2, an alternative scenario to that applied in Sections 7 and 8 is considered, where McTMA is assumed to exactly replace MIT metering. In this scenario McTMA is modeled to meter aircraft only during those periods when MIT restrictions were applied in the baseline model. This results in significantly different

metering periods than in the scenario presented in Sections 7 and 8 where McTMA metering is initiated when demand exceeds capacity. Consequently the economic benefits for this scenario are different. These benefits are presented in Table 33 below.

Table 33. Yearly increase in operating profit due to McTMA, 2003 – alternative scenario for when flights are metered.

Airport	2003 Yearly Savings [US$ / year]	
	Nominal Scenario (Metering when Demand > Capacity)	Alternative Scenario (Metering during MIT periods only)
PHL	17,275,000	12,330,000
LGA	11,810,000	2,255,000
EWR	12,817,000	6,762,000
JFK	3,289,000	2,866,000
TEB	383,000	451,000

It is clear from the results presented in Table 33 that the benefits are significantly lower under the alternative scenario where flights are metered during MIT periods only. This is because MIT metering is not applied during many of the periods when demand exceeds capacity. If restricted to meter only during period when MIT were applied, McTMA thus meters significantly fewer flights. The benefits are correspondingly lower. LGA particularly shows a very large decrease in benefits. This suggests that the greatest difference between the amount of metering under the two scenarios is at LGA. LGA operates for much of the day above capacity, so under the nominal scenario there is metering through much of the day. MIT are not, however, generally applied throughout the day, even at LGA, and thus the benefits in the alternative scenario are significantly lower.

10. Extrapolation to Future Years and Other Facilities

The McTMA assessed benefits presented in Section 8 were extrapolated to future years as described in Section 10.1 below. A procedure for extrapolating the benefits to other candidate McTMA facilities is suggested in Section 10.2.

10.1. Extrapolation to Future Years

As mentioned in Section 2.2.2, the McTMA benefits for future years were extrapolated by simulating future baseline operations without the application of McTMA and comparing them to simulated future operations with the application of McTMA. In section 4 the modeling of baseline operations is described in detail. This model was used to simulate future baseline operations, under future demand levels. The McTMA model, described in Section 5, was used to simulate future operations with the application of McTMA.

Demand Increase in Future Years

The benefit assessment for future years was calculated for demand increased according to the forecasts for each year 2010, 2015, and 2025. The FAA APO TAF [14] forecasts demand increases for these years, relative to 2003 levels, as follows:

Table 34. Forecast increases in demand over 2003 levels.

Airport	Forecast Increase in Demand from 2003 [%]		
	2010	2015	2025
PHL	25.0 %	37.8 %	63.4 %
LGA	7.7 %	7.7 %	7.7 %
EWR	28.6 %	42.5 %	70.2 %
JFK	24.4 %	41.9 %	76.7 %
TEB	14.3 %	24.5 %	44.9 %

The increase in demand at LGA is forecast to be restricted in 2006. The demand at LGA thus maintains a constant level after 2006. This level is only 7.70 % higher than 2003 demand levels. The increase in demand is not restricted at any of the other airports considered. The TAF only forecasts demand to 2020. The 2025 percentage increases are extrapolated linearly beyond 2020. A linear extrapolation is the same as used by the TAF prior to 2020.

The approach to analyze the future benefits under increased demand was to rerun the analysis under demand that is increased by randomly adding flights into the schedule.

Demand was increased randomly, but ensuring that the temporal and spatial distribution of flights through the day, and across the arrival network for each airport, was maintained. The temporal distribution of arrivals was maintained by increasing the demand in each hour by the same percentage, and distributing the new flights in each hour evenly over that hour. The spatial distribution of arrivals was maintained by randomly assigning each new flight to a network stream according to the ratios of the number of flights arriving on that stream in the baseline, to the total number of flights arriving at the airport in the baseline. The increase in schedule did not reflect the move towards increased jet to prop ratio in the mix of traffic.

MIT Increase in Future Years

In order to predict MIT in future years, the restriction generation part of the current operations baseline model described in Section 4.1, was run for future years. This model is based on ASPM demand, which is not known for future years. Since ASPM demand is derived from the actual and scheduled arrival times of flights, it is possible to estimate it using the ATAs and ETAs calculated with the increased demand of future years. In a 15-minute time period, ASPM demand is the number of flights that land in the period as well as the number of flights that were scheduled to land in or before the period, but which have not yet landed. It is also possible to estimate the scheduled demand from the ETAs alone. The ETA demand and the estimated ASPM demand were each compared to actual ASPM demand for 2003. The results of this comparison are shown in Table 35 as correlation coefficients and the square roots of the sums of the squares of the residuals, both calculated over the month of November 2003.

Table 35. Results of the comparison of actual ASPM demand with each of an estimated ASPM demand calculated from ETAs and ATAs; and an estimated scheduled demand calculated from ETAs.

Airport	Estimated ASPM demand model		ETA demand model	
	Correlation Coefficient	$\sqrt{(\sum res^2)}$	Correlation Coefficient	$\sqrt{(\sum res^2)}$
PHL	0.741	168.3	0.910	79.3
LGA	0.598	168.6	0.636	161.1
EWR	0.495	139.1	0.535	136.4
JFK	0.761	99.9	0.818	66.1
TEB	0.842	51.8	0.848	47.2

When comparing both the correlation coefficients and the square roots of the sums of the squares of the residuals, all airports show a closer correlation between actual

ASPM demand and ETA demand than between actual ASPM demand and the estimated ASPM demand. For this reason, ETA demand was taken as the best estimate of ASPM demand and hence used in the restriction generation model. The performance of this estimate of ASPM demand is shown in Table 36 by comparing the results of the restriction generation model for a period of time in November 2003 when using actual ASPM demand; and when using the ETA demand. The results are shown as fractions of the period of time for which each airport is under restriction. All days with GDP/GS are omitted from the analysis, which is also restricted to the 15 days analyzed from November 2003.

Table 36. Comparison of fraction of month restricted in November 2003, as calculated by the restriction generation model using ASPM demand as input, and using ETA demand as input.

Airport	Model using actual ASPM demand	Model using ETA demand
PHL	11.7%	6.5%
LGA	22.2%	20.1%
EWR	9.9%	2.3%
JFK	1.9%	2.4%
TEB	7.2%	5.0%

Using ETA demand under predicts at all airports except JFK where it over predicts slightly, indicating that the ETA demand estimate of ASPM demand is conservative.

Since the model is to be used to predict restrictions in future years, it is insightful to compare the restrictions predicted with actual restrictions in 2003, available from the logs. Table 37 shows the extrapolation to 2015 compared to both the model results (using ETA demand) in 2003 and the log restrictions in 2003.

At all airports, there was an increase in the fraction of the month restricted in 2015 from 2003, as expected. When comparing to the modeled restrictions in 2003, the increase is significant at both EWR and PHL due to the high increase in demand predicted by the TAF at these airports. At LGA, the increase is smaller since the demand at LGA is capped in 2006.

Table 37. The fraction of the month of November restricted according to the logs and the restriction generation model in 2003, compared to the fraction calculated by the model for 2015.

Airport	Logs	Restriction Generation Model	
	2003	2003	2015
PHL	13.7%	6.5%	20.6%
LGA	12.8%	20.1%	25.2%
EWR	6.3%	2.3%	22.4%
JFK	1.8%	2.4%	8.3%
TEB	11.7%	5.0%	12.6%

10.1.1. Extrapolation Results for Year 2015

The results presented in Section 8 were extrapolated to 2010, 2015, and 2025 by modeling the forecast increased demand at each airport and the identified increases in MIT, and comparing the McTMA simulation results to the simulated MIT baseline results. McTMA arrival delay and fuel burn were calculated on each of the 15 days, and yearly benefits calculated accordingly. These yearly benefits thus represent the yearly benefits of McTMA operations in 2010, 2015, and 2025. The assumptions made regarding the modeling are identical to those described in Section 7.2, including application of the the expected achievable increase in capacity under McTMA operations. The results are presented in Table 38 below.

Table 38. Yearly increase in operating profit due to McTMA, extrapolated to 2010, 2015, 2025 (in 2003 US Dollars).

Airport	Yearly Savings [2003 US$ / year]			
	2003	2010	2015	2025
PHL	17,275,000	56,538,000	94,779,000	363,808,000
LGA	11,810,000	28,569,000	29,480,000	29,409,000
EWR	12,817,000	96,655,000	196,430,000	520,122,000
JFK	3,289,000	14,595,000	40,253,000	172,494,000
TEB	383,000	605,000	1,635,000	12,922,000

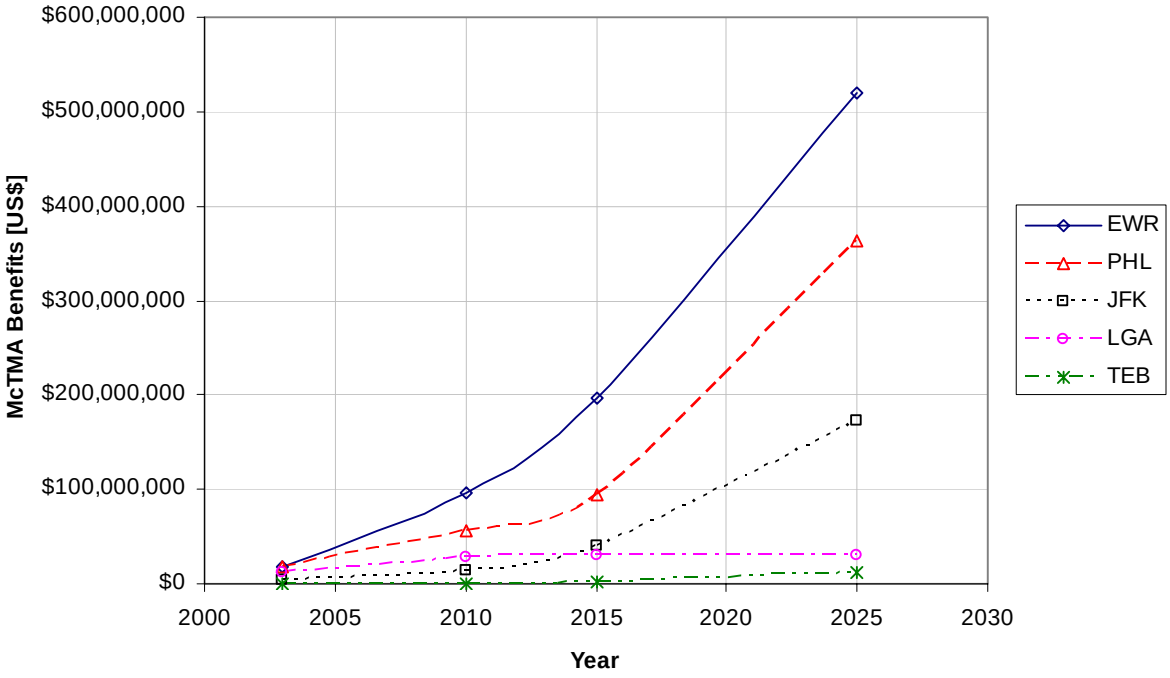

Figure 95. Increase in McTMA Benefits from 2003 to 2025.

It is important to note that the delay levels after 2003 are very high under both current operations and McTMA operations. This is expected, however, given the demand increase forecast by the FAA APO TAF [14], and the capacity applied, which was not increased above those applied for the 2003 benefits. This means that under current operations or under McTMA these levels of demand are not feasible. This observation is consistent with [15] which suggests that current demand forecasts do not adequately account for capacity constraints.

10.1.2. Remaining Issues

McTMA may be extended to include the PHL and N90 inbound flows from ZID. One approach to add the benefits from the ZID flows is to first determine how many streams the ZID flows add and to which tiers; then to assess the fraction of the benefits associated with similar streams that were already included in the analysis; and finally to augment the benefits by this fraction. However, the impact of ZID's flows into N90 and PHL was not addressed in the extrapolation of the benefits to future years. This makes the current assessment conservative, as it does not include the added benefit from the ZID flows.

10.2. Extrapolation to Other Facilities

SAIC identified a number of potential McTMA sites in task order 54 [16]. Their criteria included: facility operational requirements, ARTCC/TRACON boundaries, airspace complexity, airport capacity/delay, and weather phenomena. Metrics that represent these criteria may be used to scale and extrapolate the McTMA benefits identified for PHL and N90 to other facilities. However, as [16] concluded, common metrics do not currently exist for most of these criteria, such as airspace complexity and weather phenomena. Developing such metrics would require extensive research beyond the scope of this task. Some metrics can, however, be developed to represent one of these criteria, i.e. airport capacity, which was identified as the main parameter impacting McTMA benefits.

Therefore, the method suggested for the extrapolation of the PHL and N90 McTMA benefits to other facilities is based on metrics representing airport capacity and the airport demand relative to this capacity. Three metrics for demand and capacity were identified. These metrics were selected such that they can be easily computed using ASPM data. The metrics are as follows:

1. The total number of arrival operations at an airport per day (directly representing demand),
2. The percentage of the year for which arrival demand exceeds reported arrival capacity (representing when McTMA will be operated, as it is to be operated when demand exceeds capacity), and
3. The utilization of the airport (representing how close to available capacity the airport operates).

Utilization was estimated as the ratio of the average actual throughput to the 99th percentile of the actual throughput at the airport.

An equation relating these three metrics to the benefits was derived from the following relation from queuing theory (see Section 2.2.1 equation (2)):

$$Delay \propto \frac{\lambda}{1-\lambda/\mu}$$

where λ represents demand and μ represents capacity.

Benefits are considered proportional to delay. This was observed in the extrapolation to future years. An exponential increase in delay led to an exponential increase in benefits, as shown in Section 10.1.1.

The physical behavior of the queuing relation may be considered by recognizing that the arrival demand λ is equivalent to the average number of arrival operations (*ArrOps*). When arrival operations tend to zero, the delay and hence benefits will also tend to zero. Similarly, when demand tends to the real capacity, the delay, and hence benefits, tend to infinity.

λ/μ is approximated by the product of the percentage of the year for which arrival demand exceeds reported arrival capacity (this is a measure of the duration of metering)

and the utilization of the airport (this is a measure of the average difference between demand and capacity):

$$P_{D>C} \times Util \propto \lambda / \mu$$

where $P_{D>C}$ is the percentage (between 0 and 100) of the year for which arrival demand exceeds reported arrival capacity (measured per 15 minute interval), *Util* is the utilization of the airport.

This yields the following equation relating the three metrics to the benefits:

$$Benefits = A\left[\frac{ArrOps_{perday}}{1 - B(P_{D>C})(Util)}\right]$$

where A, and B are constants.

The least squares fit of this equation to the benefits quoted in Table 26 for the assumption of an increase in capacity, and the three metrics yielded a correlation coefficient of 0.966 and a square root of the sum of the squares of the residuals of $6,246,000.

The resulting constants resulted in the following equation:

$$Benefits = 19{,}366 \times \left[\frac{ArrOps_{perday}}{1 - (0.0213)(P_{D>C})(Util)}\right] \quad (13)$$

The data used for the metrics in this analysis can be easily extracted from the ASPM database (http://www.apo.data.faa.gov/faamatsall.HTM). Arrival demand comes from the ARRDEMAND field in the database, reported arrival capacity from the AAR field, and number of arrival operations (or throughput) from the EFFARR field. Utilization is calculated from the average and 99th percentile of this field over a year.

As an example the above model was used to predict the benefits at LGA. According to the ASPM database the inputs were as follows:

1) The percentage of the year for which arrival demand exceeded reported arrival capacity, $P_{D>C}$, was 37.73 %.
2) The average number of arrival operations per day over the year, $ArrOps_{per\ day}$, was 510.6 aircraft.
3) The utilization, calculated as the ratio of the average actual throughput to the 99th percentile of the actual throughput at the airport was 0.51.

The resulting McTMA benefits calculated per year were: US $ 16,791,000. This compares to the results presented in Table 26 of $15,535,000. The error is 8.1%.

11. Conclusions and Recommendations

The benefit estimates of McTMA assessed in this study are believed to be realistic and robust for a number of reasons:

1. A large sample of days was analyzed – fifteen days of November 2003. These days represented a random and wide range of metering conditions (demand exceeding capacity) both in terms of duration and severity.
2. Days or periods of time when the system was thought to be restricted by constraints irrelevant to McTMA (or not included in this study) such as local restrictions not related to runway capacity or strategic restrictions like ground delay programs, were excluded. This was done in order to limit the benefits assessment conservatively to those delays that McTMA is believed to mitigate.
3. The benefit estimates are also believed to be realistic because they resulted from comparing a McTMA model of time based metering to a model of baseline operations using distance based metering. By using modeled baseline operations care was taken to model only procedures and dynamics that are relevant to McTMA – situations of demand exceeding capacity that require metering through distance based metering (as opposed to rerouting or GDP for example).
4. The delays were measured with respect to unimpeded estimated times of arrival. These unimpeded times were derived from statistical models based on analysis of historical track data. The statistical unimpeded times took wind, aircraft type, and runway configuration into account. These statistical models were compared to estimated times of arrival computed from trajectory synthesis (the CTAS Trajectory Synthesizer process) based on flight plan, wind, and aircraft performance. The statistical models compared well to the trajectory synthesis models as reported in Section 3.2.

The benefit estimates of McTMA assessed in this study are also believed to be conservative due to a number of reasons:

1. The benefits assessment focused on a subset of McTMA functions as described in Section 4. These functions included time based metering by generating STAs, delay feedback and capacity distribution, dynamic metering, tiered metering, coordination between multiple facilities, demand visualization, and scheduling of internal departures. Functions such as runway assignment, which are planned but not currently implemented in McTMA, or other functions envisioned in the future, were not analyzed. These additional functions may be analyzed in future research and extension of the benefits assessment.
2. Benefit mechanisms were derived for each of the McTMA functions and described in charts in Section 6. However, due to time and resource limitations, not all of the benefit mechanisms were modeled and analyzed in this study. The benefit mechanisms not analyzed are indicated with dashed arrows in the benefit mechanisms charts in Section 6. Most notably these include:

- Switching flights between arrival flows either to avoid severe weather impacting an arrival gate or to more optimally offload aircraft to other facilities in the face of congestion.
- Improved decisions to shut off the traffic in extreme situations causing no-notice holding or gridlock.
- Interaction with facilities outside of the McTMA system such as the Command Center or other ARTCCs. This interaction may lead to reduction in the use or severity of ground delay programs or MIT restrictions upstream of the McTMA system. Demand into the McTMA system would then increase to take advantage of the increased throughput and utilization due to time based metering.
- Impact on the takeoff time of internal departures. In this study internal departures were assumed to takeoff at their actual takeoff time, and additional delay was absorbed on the ground if needed. McTMA may impact this takeoff time resulting in more optimal distribution of delay between air and ground.

The analysis of these additional benefit mechanisms may result in additional benefits assessment for McTMA. Additional benefit mechanisms that may have been missed in this study or for additional McTMA functionality may also be derived in the future.

3. The benefit estimates also depended on a number of assumptions made about the McTMA operation in the field and about corresponding modeling parameters. Care was taken to the extent possible to make conservative assumptions and to consult NASA's McTMA researchers and their experience with the tool and in the field. Through sensitivity analysis a range was tested for many of these assumptions and parameters to provide a range of corresponding benefit estimates and to assess how much of an impact such assumptions and parameters have on the benefit estimates. Most notably, the benefit estimates depended largely on the assumption about the capacity limits imposed by McTMA on the system. A range of such capacities were tested, ranging from imposing the same capacity applied to the baseline to imposing the maximum capacity believed acceptable based on safety and operational concerns. Also while it was assumed that McTMA's time based metering would be applied during times of demand exceeding capacity even if MIT was not applied, a more conservative scenario was provided in the sensitivity analysis in which time based metering only replaced MIT during the times that MIT was applied.
4. Also some of the benefit mechanisms that were modeled and analyzed were only captured partially due to simplifying assumptions made. For example, the effect of the dynamic nature of metering (updating STAs every 12 seconds) was not modeled explicitly. Rather it was captured partially by adjusting for the errors in meeting the STAs from tier to tier. A number of such simplifications needed to be made and care was taken to make them conservatively.

Future Work and Extensions

Additional research may be pursued based on the insights gained in this study. Some of these research areas are outlined below:

1. Future research may identify the causes of delay and inefficiency in the actual traffic data. By listing the sources of delay in the actual data baseline it should be possible to identify those sources that McTMA is expected to mitigate and those that it is not. This analysis requires a detailed diagnosis of the delays and their causes. Each aircraft should be classified by associating a number of characteristics to it, such as the restrictions that affected it during its flight, the number of aircraft it encountered, its speed, and other factors. More detailed facility logs may be required to obtain more detailed information.
2. Future research may also focus on identifying more accurate assumptions about the operation of McTMA in the field. These assumptions are key factors in deciding on the delay sources that McTMA may mitigate. These assumptions regarding the operation of McTMA should be investigated in future research based on more consultation with McTMA researchers and their experience with McTMA at PHL.
3. Future research may continue the effort to build a model of current operations. This requires further expert elicitation and more detailed facility logs and track data for extended periods of time.
4. While the demand was modeled with high fidelity in this study – using detailed arrival flow networks, carefully selected meter fixes, and Host track data with 12 second resolution to estimate unimpeded transition times – there is still need to further validate and verify the models. Future research should investigate alternative statistical models and flight path based models to validate the models developed in this study. More refined classification of flights may also increase the accuracy of the unimpeded transition times.
5. Capacity models were generated in this study using ASPM data, which are counts of traffic in 15-minute periods. While the models were insightful and sufficient to identify capacity limits, they may be refined and confirmed using more accurate Host track data. Also more research is warranted to further investigate the TRACON as a main flow constraint, which was observed in this study.
6. The five airports analyzed in this study provide a range of different characteristics. The benefit estimates were different at the different airport, but minimal time was available to conduct comparative analysis and gain insight on the causes of the differences. Future research may conduct more such comparisons, which may also lead to understanding of the differences in McTMA operation at the different airports. For example, the hypotheses made in this reports about the PHL benefits due to increased capacity being higher than the other airports' benefits because of the bank schedule may be investigated further and confirmed through studies at other airports.
7. In this study days with ground delay program were excluded neglecting the impact of McTMA in terms of reducing GDP delays. The interaction between time-based metering and strategic ground delay may be investigated in future research. In

addition the impact on MIT restrictions upstream of the McTMA system and on internal departures delay on the ground may also be investigated.

8. Future research may also investigate determining delay absorption limits for the TRACON and airspace sectors at optimal values beyond which no further gains in throughput may be achieved. This may build on the analysis of the throughput drop off as demand increases.
9. Future research may investigate the drop off in arrival throughput at high levels of demand. This may be investigated at additional airports not analyzed in this study in order to confirm the behavior as a general characteristic. The drop off should also be confirmed using models based on more accurate Host track data for demand and throughput measurement. Causes of the drop off should be investigated by combining other sources of information about restrictions for example.
10. Extensions to the concept of McTMA may be researched in the future based on insights gained in this study. One such extension is the addition of advisories from rerouting around convective weather within the boundary of the McTMA system. This may be achieved, for example, by reallocating aircraft to different stream, which is not currently a function of McTMA. With the extension of time based metering to three tiers, the accommodation of local weather constraints within the system becomes an important aspect. These local constraints may not be reflected by the acceptance rate, which is the main explicit constraint that TMA and McTMA consider currently.
11. Future research may also be conducted to develop a more comprehensive economic model for airline response and collaboration.
12. Future research may also investigate the benefits of McTMA for future years applying improved and more realistic demand forecasting that takes system constraints into account.

References

[1] Swenson, H. N., T. Hoang, S. Engelland, D. Vincent, T. Sanders, B. Sanford, and K. Heere, *Design and Operational Evaluation of the Traffic Management Advisor at the Fort Worth Air Route Traffic Control Center,* 1st USA/Europe Air Traffic Management R&D Seminar, Saclay, France, June, 1997.

[2] Farley, T. C., J. D. Foster, T. Hoang, K. K. Lee, *A Time-Based Approach to Metering Arrival Traffic to Philadelphia,* AIAA-2001-5241, First AIAA Aircraft Technology, Integration, and Operations Forum, Los Angeles, California, October 16-18, 2001.

[3] Hoang, T., T. Farley, J. Foster, and T. Davis, *The Multi-Center TMA System Architecture and Its Impact on Inter-Facility Collaboration,* AIAA Aviation Technology, Integration and Operations (ATIO) Conference, Los Angeles, CA, October 1-3, 2002.

[4] Wong, G. L., *The Dynamic Planner: The Sequencer, Scheduler, and Runway Allocator for Air Traffic Control Automation* NASA/TM-2000-209586, April 2000.

[5] Landry, S., T. Farley, J. Foster, S. Green, T. Hoang, and G. Wong, *Distributed Scheduling Architecture for Multi-Center Time-Based metering*, AIAA-2003-6758, 3rd Annual Aviation Technology, Integration and Operations (ATIO) Conference, Denver, Colorado, November 17-19, 2003.

[6] Couluris, G. J., Weidner, T. J., *Terminal Airspace Decision Support Tool Functions Preliminary Potential Benefits* Analysis, Final report Research Task Order 10.

[7] Metron Aviation, *Multifacility TMA: Benefits Analysis for Philadelphia Installation*, Final report Research Task Order 33.

[8] Metron Aviation, *Regional Metering: Benefits Analysis for Regional Metering*, Final report Research Task Order 71.

[9] William Chan, NASA CTAS Engineering, May 2003.

[10] FAA-APO-98-8, *Economic Values for Evaluation of Federal Aviation Administration Investment and Regulatory Programs*, FAA Office of Aviation Policy and Plans, June 1998, http://api.hq.faa.gov/economic/toc.htm.

[11] DOL BLS, *Consumer Price Indexes, Inflation Calculator*, US Department of Labor, Bureau of Labor Statistics, July 2003, http://www.bls.gov/cpi/

[12] BTS OAI, *Air Traffic Statistics and Airline Financial Statistics*, Bureau of Transportation Statistics, Office of Airline Information, June 2003, http://www.bts.gov/oai/indicators/top.html#DomOpFinance.

[13] BTS OAI, *Quarterly Operating Financials: Domestic/International/System*, Bureau of Transportation Statistics Office of Airline Information, June 2003, http://www.bts.gov/oai/indicators/qtropfinan.html

[14] FAA APO, *Terminal Area Forecast*, Federal Aviation Authority, Aviation Policy and Plans, June 2003, http://www.apo.data.faa.gov.

[15] Kostiuk P., Lee D., Long D., *Closed Loop Forecasting of Air Traffic Demand and Delay*, 3rd USA/Europe Air Traffic Management R&D Seminar, Napoli, June 2000.

[16] SAIC, *Investigation of Implementation Sites for Multi-Center Traffic Management Advisor (McTMA) AATT Decision Support Tool*, Task Order 54 Final Report, December 2000.

Appendix A: PHL and N90 Arrival Flows

A.1 Philadelphia TRACON (PHL) Flows

PHL has 4 primary arrival fixes: MAZIE and BUNTS, from ZNY; and TERRI and Cedar Lake (VCN) from ZDC. These arrival fixes are shown in black in Figure A.1. The arrival flows into Philadelphia TRACON (PHL) are also presented in Figure A.1.

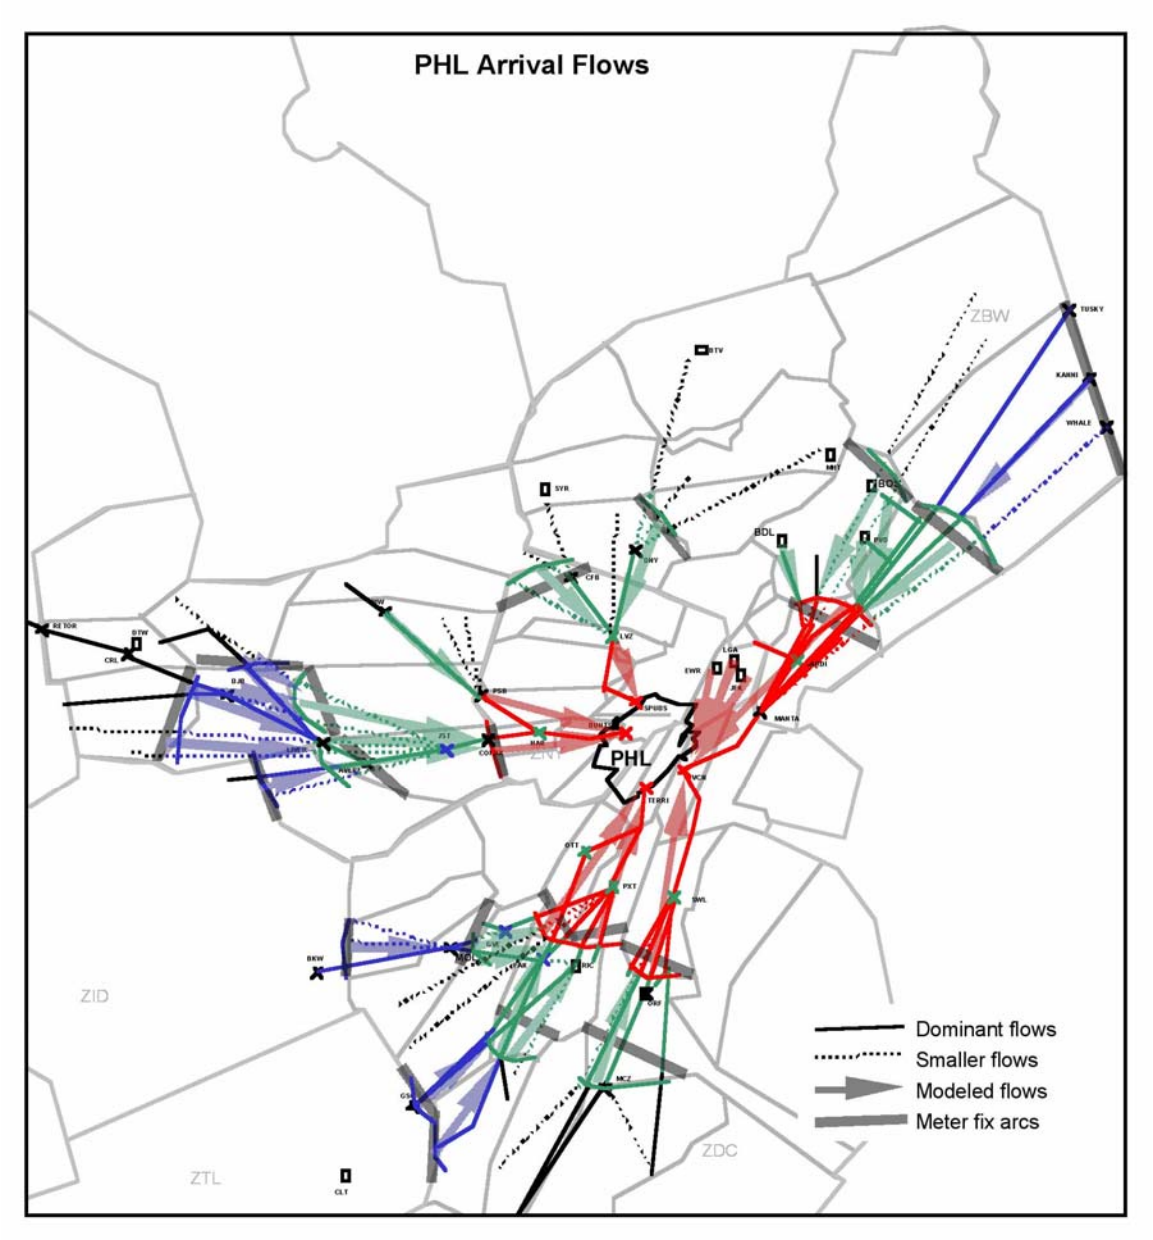

Figure A.1. Air traffic flows into Philadelphia TRACON.

Arrival traffic to PHL is to be metered at three tiers, indicated on Figure A.1 by red arrows (tier1), green arrows (tier 2), and blue arrows (tier 3). Each arrow represents a different flow that was modeled. Note that not every single actual flow was modeled. Instead flows were grouped when the flows had similar length and location. Similar length and location means that the transition times for the flows are likely to be very similar, and can thus be modeled as a single flow.

Metering arcs are used to a greater extent in the above flow network than is expected to be used operationally. This is to ensure that as much traffic as possible is captured, as aircraft routes are identified in the model according to tracks, and not flight plans, as in TMA. Meter arcs are illustrated on the figure by thick gray lines, with colored arcs associated with them. The gray lines show the metering windows through which the aircraft must pass to be considered to be on that flow, while the arcs are where crossing times and locations are recorded.

McTMA deployment for PHL is currently underway, and the above fixes have only been tentatively established, and will be adjusted based on feedback from operational testing.

A.2 New York TRACON (N90) Flows

Based on feedback from McTMA NASA researchers it is assumed that McTMA will only be implemented for Newark, LaGuardia, JFK and Teterboro airports. Each of Newark, LaGuardia and JFK have separate arrival fixes. These are presented in Table 39 below, and were determined according to the Standard Operating Procedures (SOP) for N90, letters of agreement (LOAs) between N90 and the appropriate ARTCC, and interviews with TMCs.

Table 39

Destination Airport	Origin ARTCC	Arrival Fix	Traffic
EWR	ZBW	SHAFF	Jets/Hptp
		LEMOR	Props
	ZNY	PENNS	Jets/Hptp
	ZDC	Yardley (ARD)	TBD
LGA	ZBW	VALRE	Jets/Hptp
		NOBBI	Props
	ZNY	LIZZI	Jets/Props
	ZDC	Robbinsville (RBV)	TBD
JFK	ZBW	Calverton (CCC)	All
		LOVES	Props
	ZNY	LENDY	All
		CAMRN	TBD
		ZIGGI	Props
TEB	ZBW	LEMOR	Jets/Props
	ZNY	MUGZY	Jets/Hptp

In Table 39, Hptp refers to high performance turboprop aircraft, and TBD (to be determined) to information that is unclear and must be determined through further

interviews with TMCs. Particularly, Yardley (ARD) and Robbinsville (RBV), the primary arrival fixes for Newark and LaGuardia from ZDC, were flipped in the summer of 2002, which has improved operations significantly, as the two flows do not need to be crossed, as before. The details of the traffic on each fix are still to be determined from TMCs at N90.

The primary flows into New York TRACON have been identified, and are presented in Figure A.2 to Figure A.5.

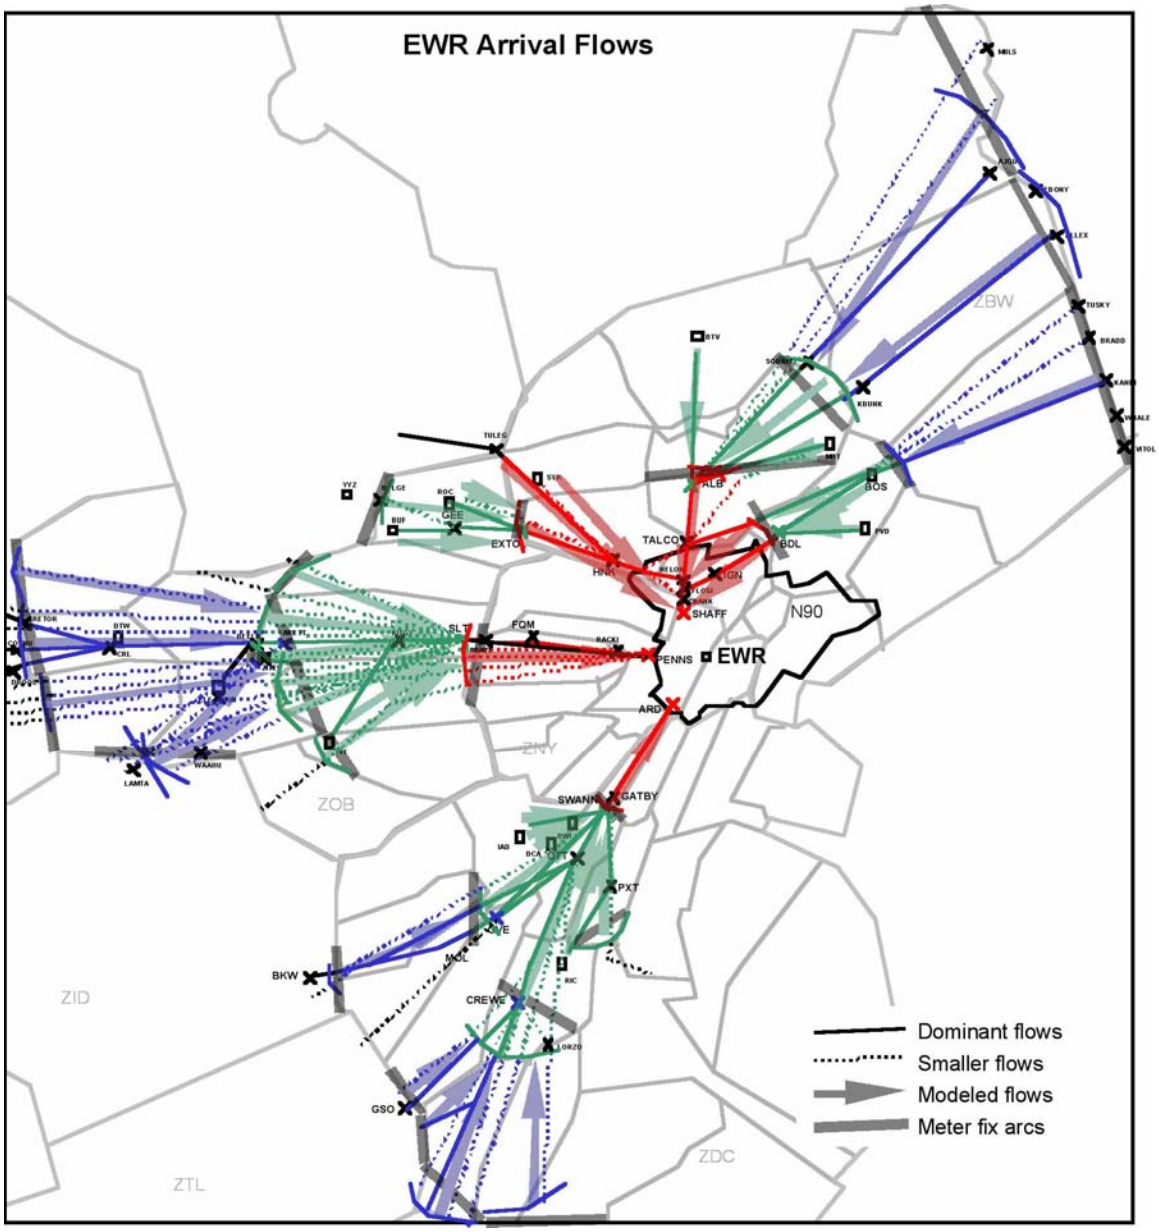

Figure A.2. Air traffic flows for Newark International Airport (EWR).

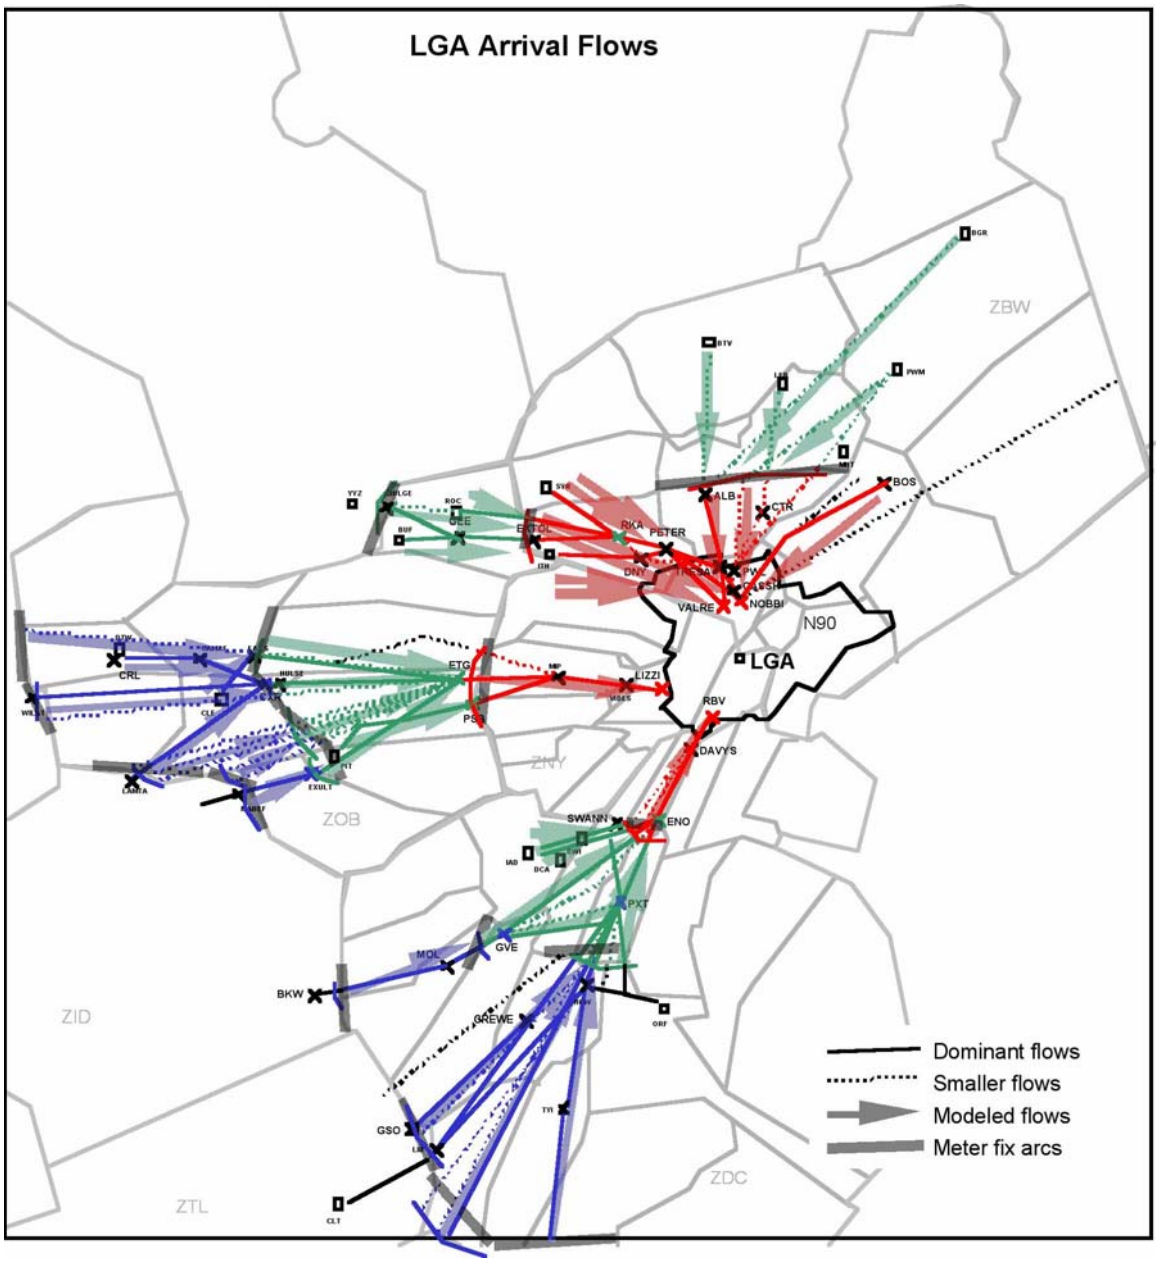

Figure A.3. Air traffic flows for LaGuardia Airport (LGA).

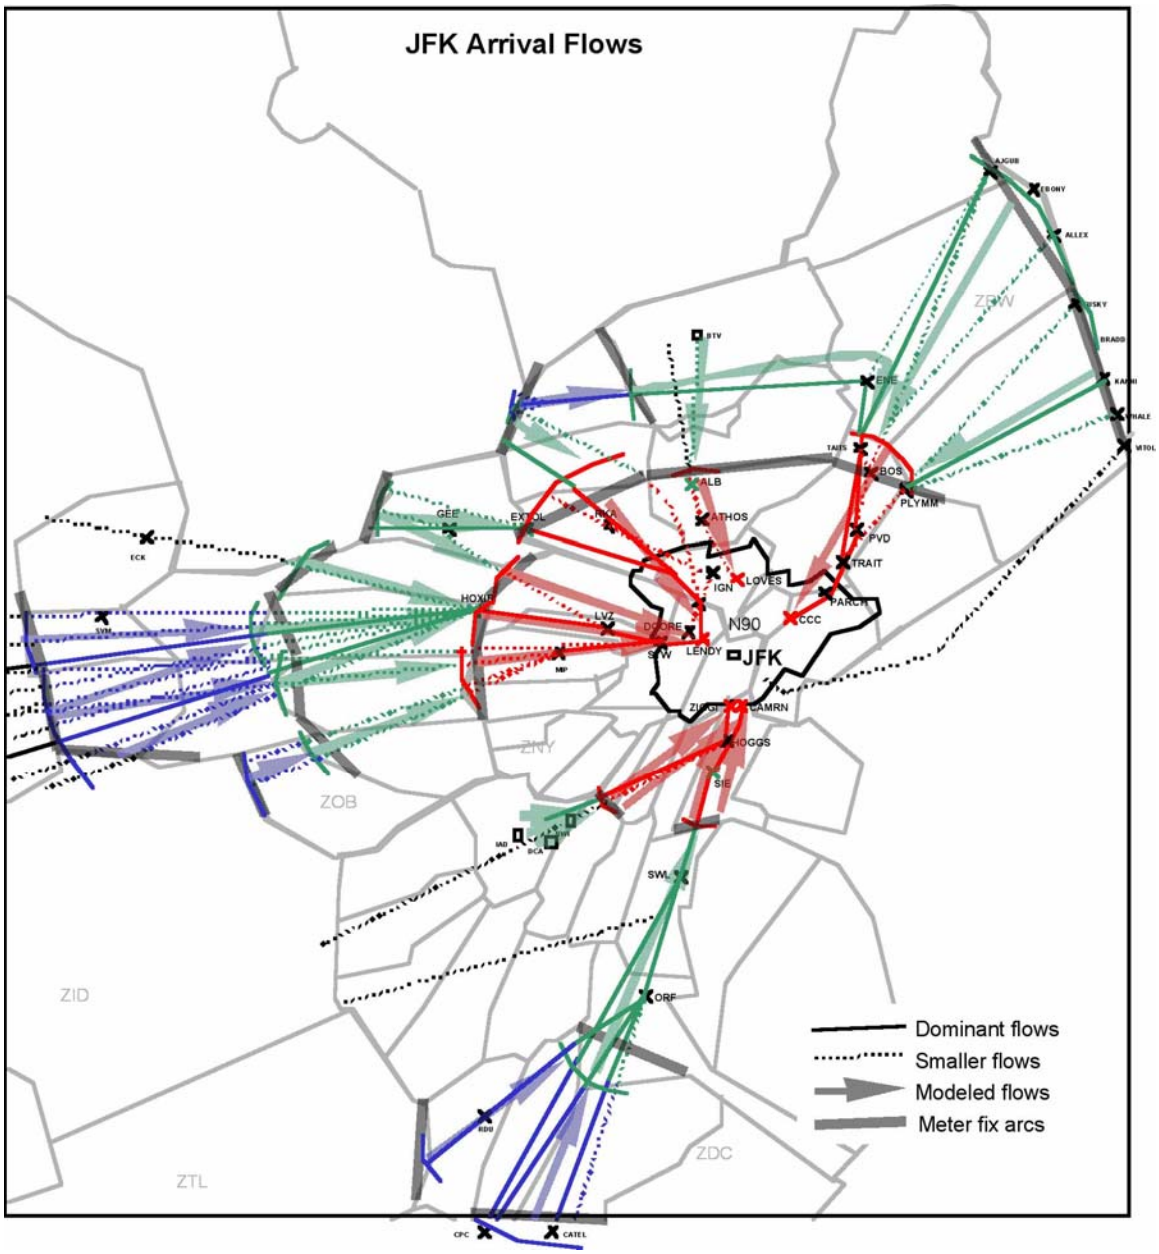

Figure A.4. Air traffic flows for John F Kennedy International Airport (JFK).

Flows into JFK from ZDC must transition from ZDC to ZNY, and then to N90. There is no direct transition from ZDC to N90 for JFK traffic.

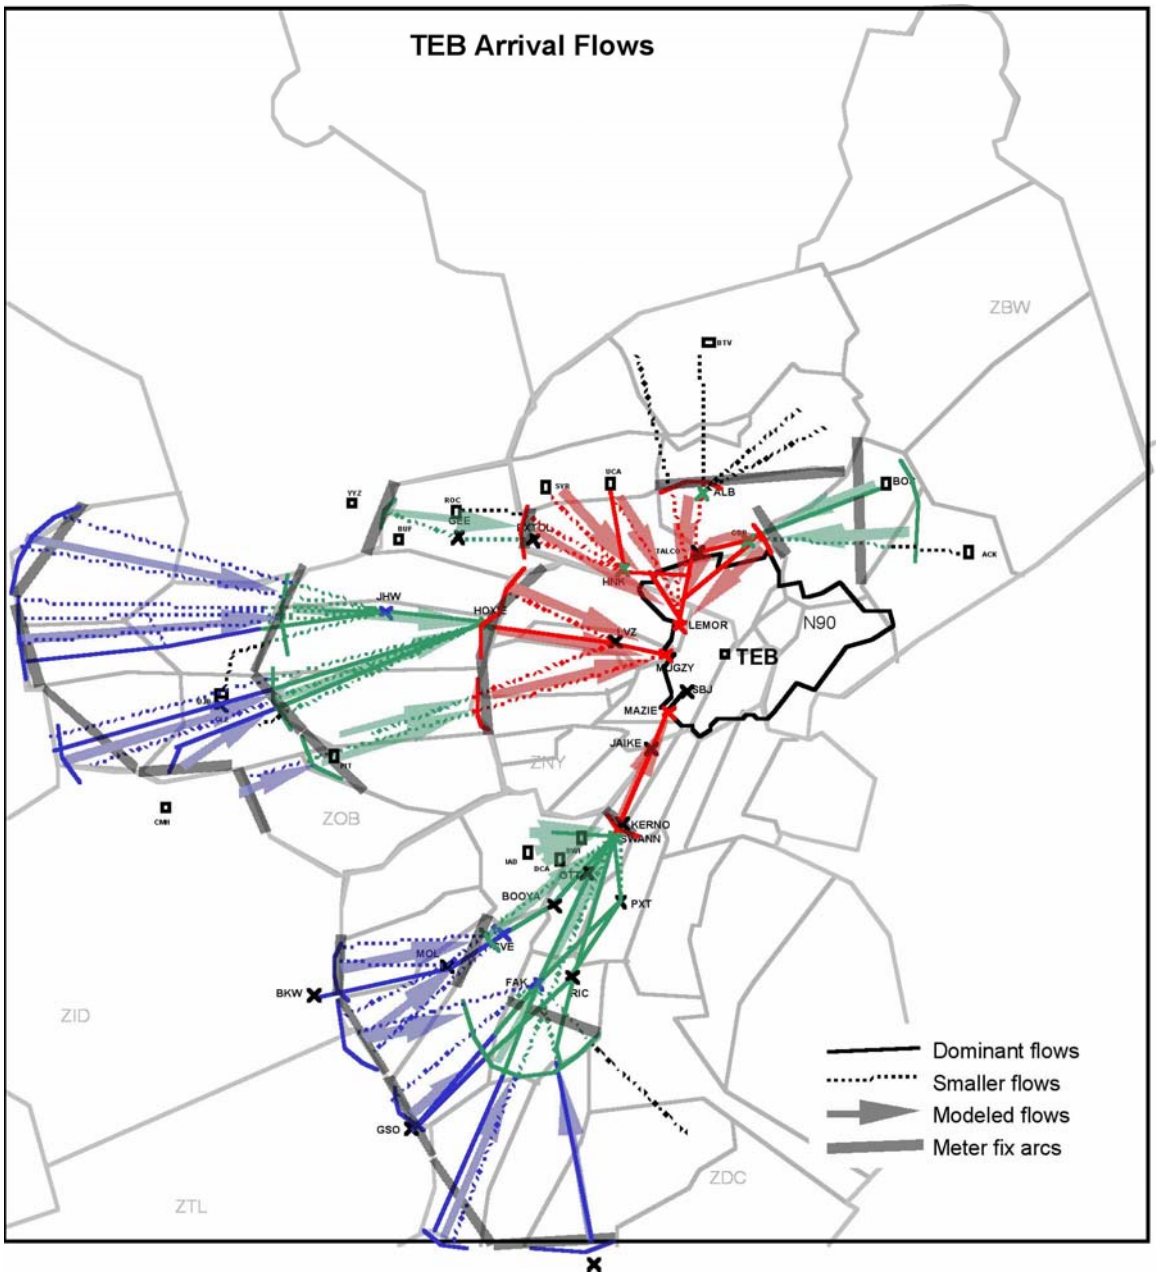

Figure A.5. Air traffic flows for Teterboro Airport (TEB).

The flows presented in Figure A.2 to Figure A.5 were identified according to STARS, flows illustrated in presentations by ZOB (*Traffic Management Unit*, by Cleveland ARTCC, FAA) and ZBW (*Boston ARTCC*, by Boston ARTCC, FAA), and flows presented by NASA for PHL (*Distributed Scheduling for McTMA*, by Todd Farley and Steve Landry, NASA Ames Research Center), and according to host track data from September 12, September 17, and September 19, 2002.

Appendix B: ASPM Capacity Analysis

Following is a capacity analysis of the four primary airports in New York (EWR, JFK, LGA and TEB), and of Philadelphia airport (PHL). The analysis is performed using the Aviation System Performance Metrics (ASPM) database, obtained through NASA and the FAA (http://www.apo.data.faa.gov/faamatsall.HTM). The results of the capacity analysis are separated by airport and airport configuration. Because the data used in the simulation of McTMA is from August 2002 and September 2002, the frequency of the different configurations at each airport is plotted for these two months. The demand at the airports was not, however, as great in these months as it was before September 11, 2001. In order to accurately analyze the capacity of the airports the data analyzed must include the periods of highest demand. Consequently, for the capacity analysis, data from January to August of 2001 was chosen for processing. The summer months are the period of highest demand, so data from the summer of 2000 was also included to enhance the data from the summer of 2001. This choice of data ensures that the maximum airport capacities determined in the analysis are representative of the actual maximum capacities of the airports.

B1. Newark International Airport EWR

AAR Frequency

The frequency of different AARs are plotted below. The data used is that from the summer of 2000, and from the winter, spring, and summer of 2001. The data used is per hour, as AARs are reported as hourly rates.

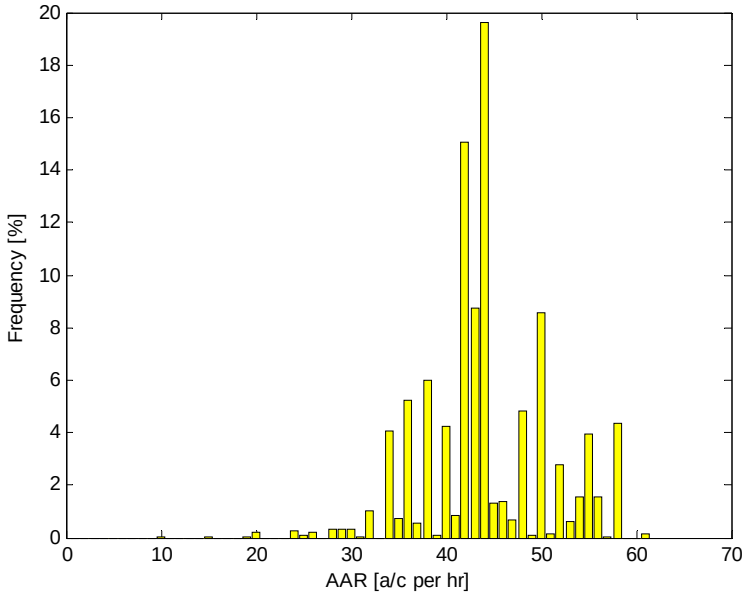

Figure B1. Frequency of reported AARs.

AARs of 34, 38, 40 and 44 are examined in more detail below. These are the AARs reported on the days analyzed in the RTO 77 study.

B1.1. Airport Acceptance Rate of 34 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below. The data used is that from the summer of 2000, and from the winter, spring, and summer of 2001.

In the first of the charts below, quarter hourly throughput is plotted against corresponding demand. In order to represent rates by the hour the values for throughput and demand are multiplied by four. All data points are plotted. This allows for the identification of the absolute maximum arrival capacity. Quarter hourly intervals may include some intervals with only arrivals, and may include intervals with consistently high demand. The capacity identified would thus not be sustainable over time, but is applicable for short periods.

The frequency of the data is presented alongside the raw data. The gray scale denotes the number of data counts in 4 aircraft per hour by 4 aircraft per hour bins. The maximum frequency on the scale to the right is approximately half the highest number of data points in any single bin.

The frequency of the demand is presented beneath the raw data. This chart allows one to identify if there are enough data points at critical points (such as the drop-off in throughput) to verify the conclusions.

Finally, alongside the frequency of demand, the moving average of the throughput is plotted against demand, with error bars representing one standard deviation in each direction. The window size for the calculation of the moving average is 10 aircraft per hour. The point at which the data drops off is identified heuristically, as that point at which the average throughput drops off at 2% over one bin (4 aircraft per hour), and for which the average of the rest of the points after this bin is lower than the 80% of the value of the average throughput in this bin. This heuristic was identified after manual experimentation with the data. A hyperbolic fit to the moving average for those points, before any drop-off, is also plotted on the graph. The hyperbolic fit asymptotes to the 45° line on the left (throughput equal to demand) and to the horizontal on the right, which represents the maximum throughput according to the trend of the moving averages before the drop-off.

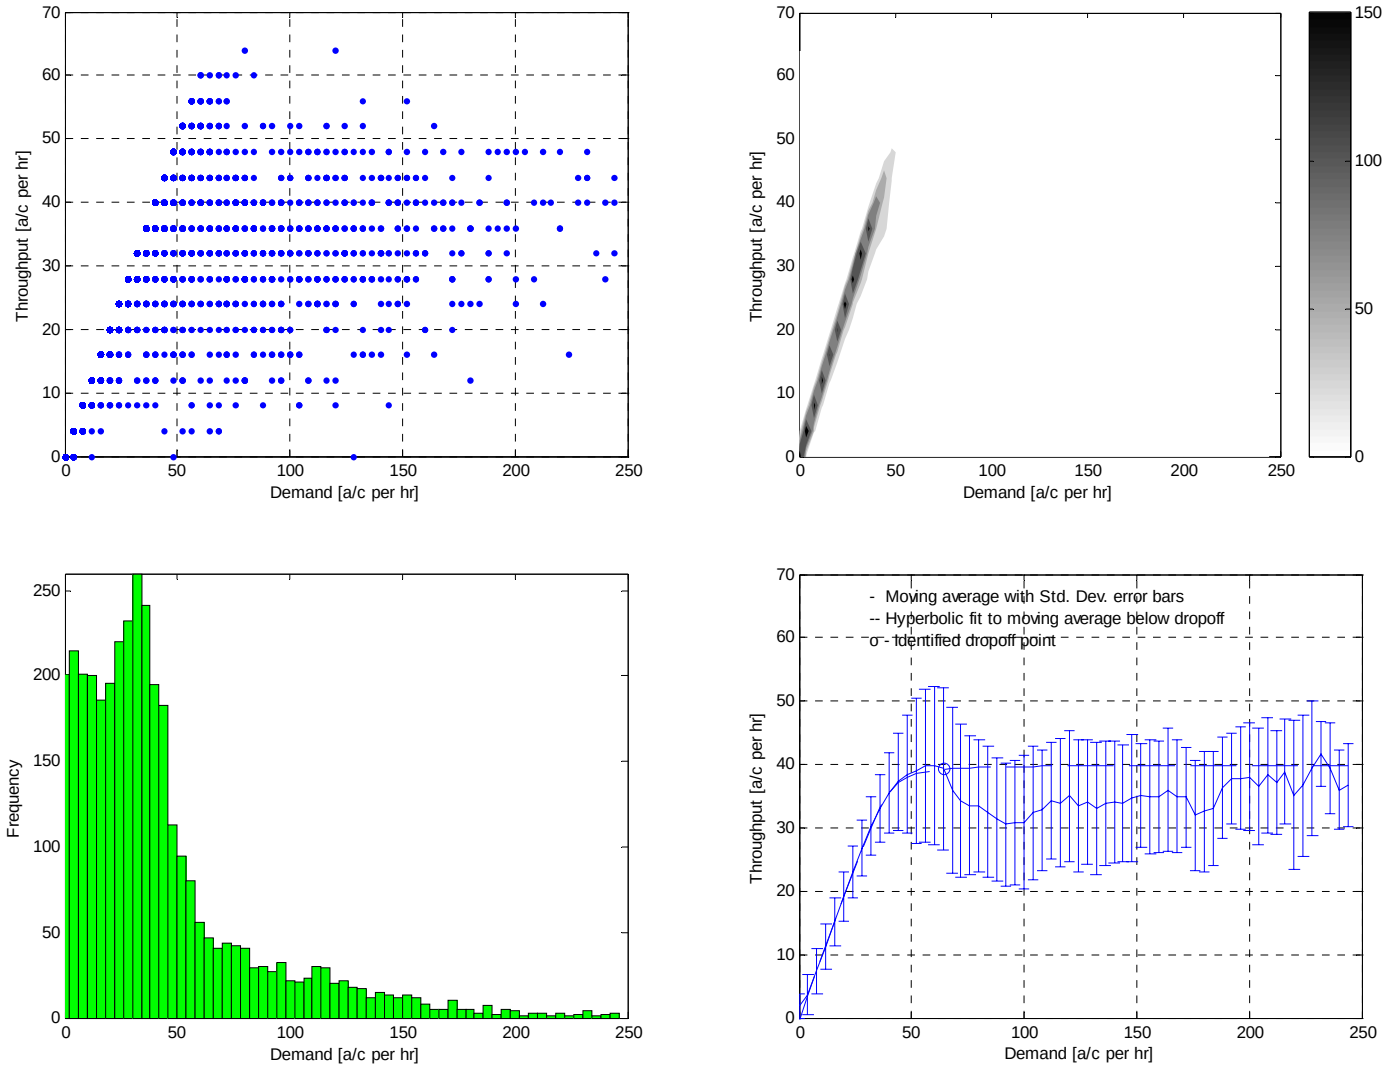

Figure B2. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

According to these graphs, the 98, 99 and 100th percentile of the throughput is calculated, along with the lowest corresponding demand. These values are presented in the table below. These values of throughput represent the capacity of the airport, while the corresponding values of demand represent the limits of demand that are high enough to reach the capacity of the airport. Also presented in the table below are the variables associated with hyperbolic curve fit (the asymptote and vertex 'a'). The asymptote also represents the capacity of the airport according to the averages of the data. Finally the maximum value of the moving average, to the left of the drop-off point; and the coordinates of the drop-off point, are presented.

Percentile of Throughput	Demand [a/c per hr]	Throughput [a/c per hr]
98	52	52
99	56	56
100	80	64
Hyperbolic Fit Asymptote		40
Hyperbolic Vertex 'a'		4
Moving Average Max	56	40
Moving Average Drop-off	64	39

B1.2. Airport Acceptance Rate of 38 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

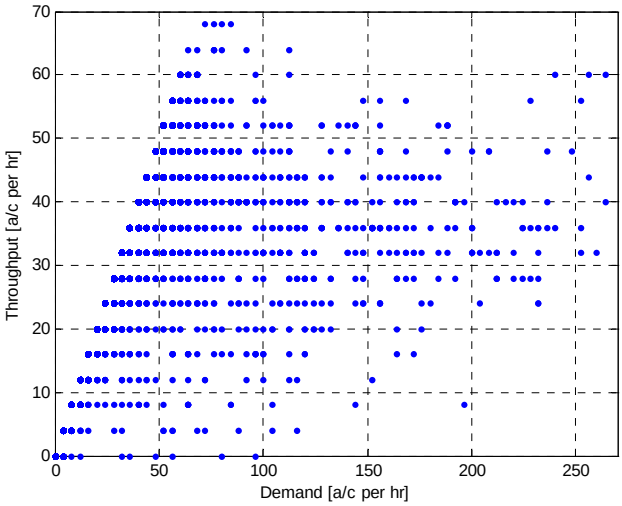

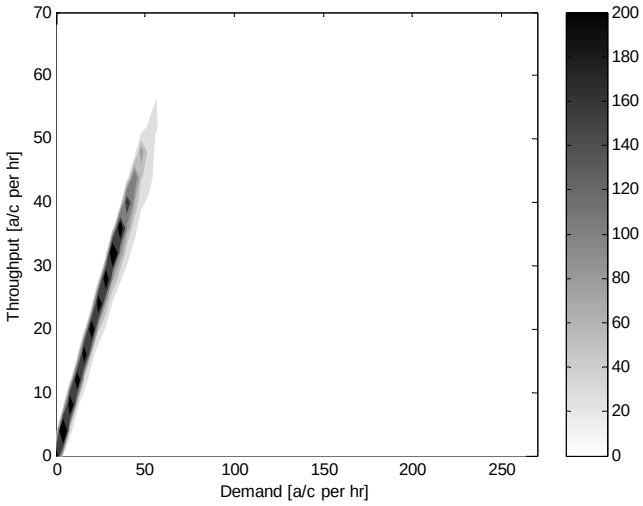

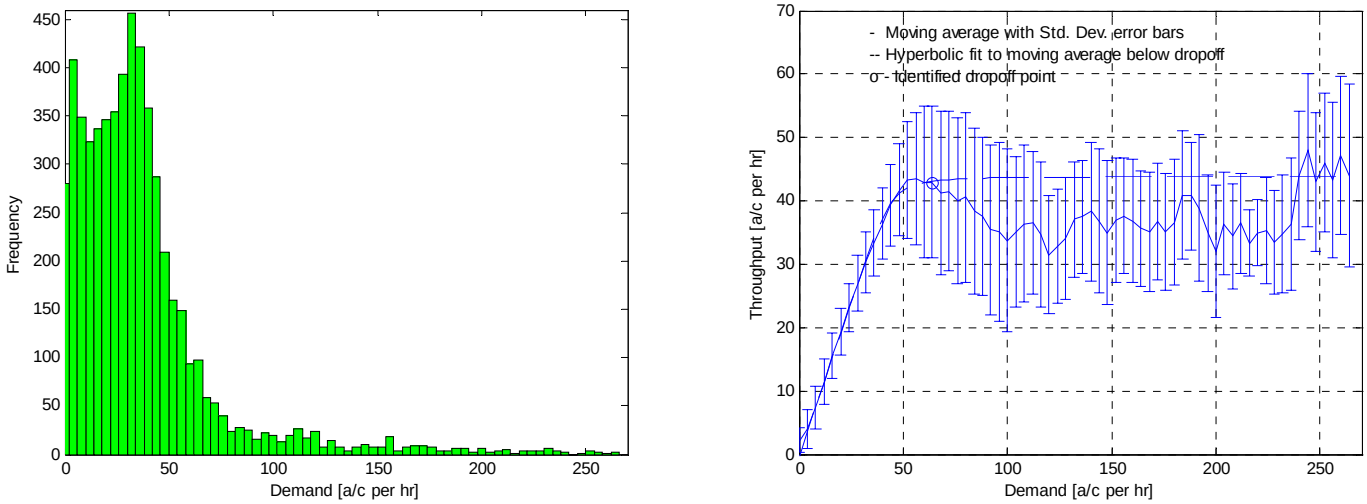

Figure B3. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	52	52
99	56	56
100	72	68
Hyperbolic Fit Asymptote		44
Hyperbolic Vertex 'a'		4
Moving Average Max	56	44
Moving Average Drop-off	64	43

B1.3. Airport Acceptance Rate of 40 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

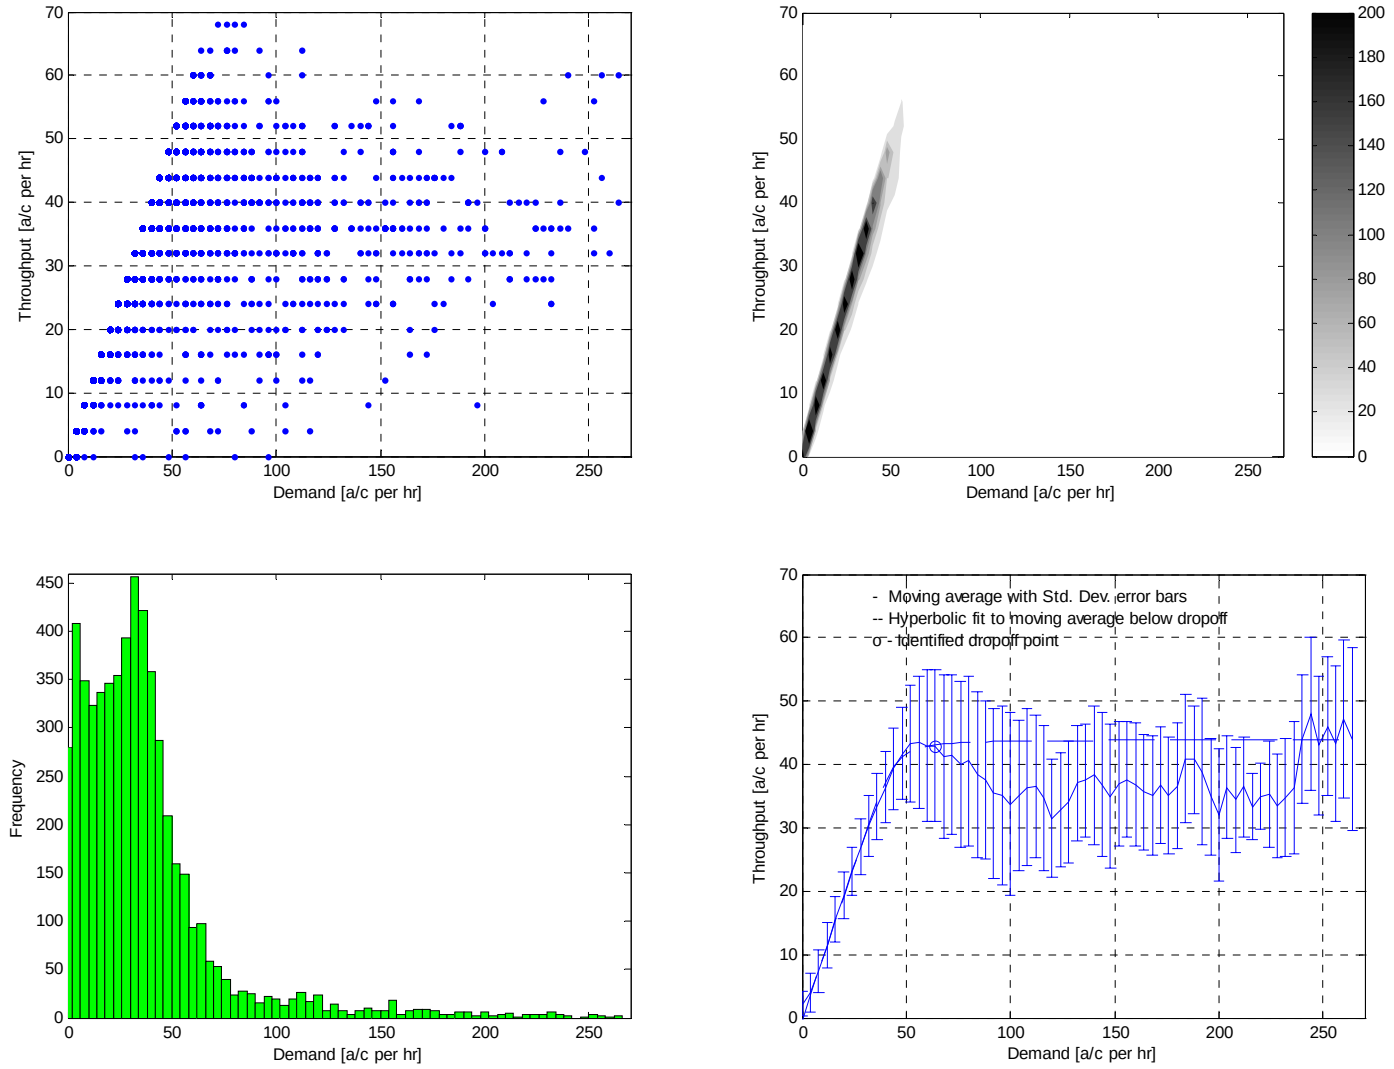

Figure B4. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	52	52
99	56	56
100	72	68
Hyperbolic Fit Asymptote		44
Hyperbolic Vertex 'a'		4
Moving Average Max	56	48
Moving Average Drop-off	64	48

B1.4. Airport Acceptance Rate of 44 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

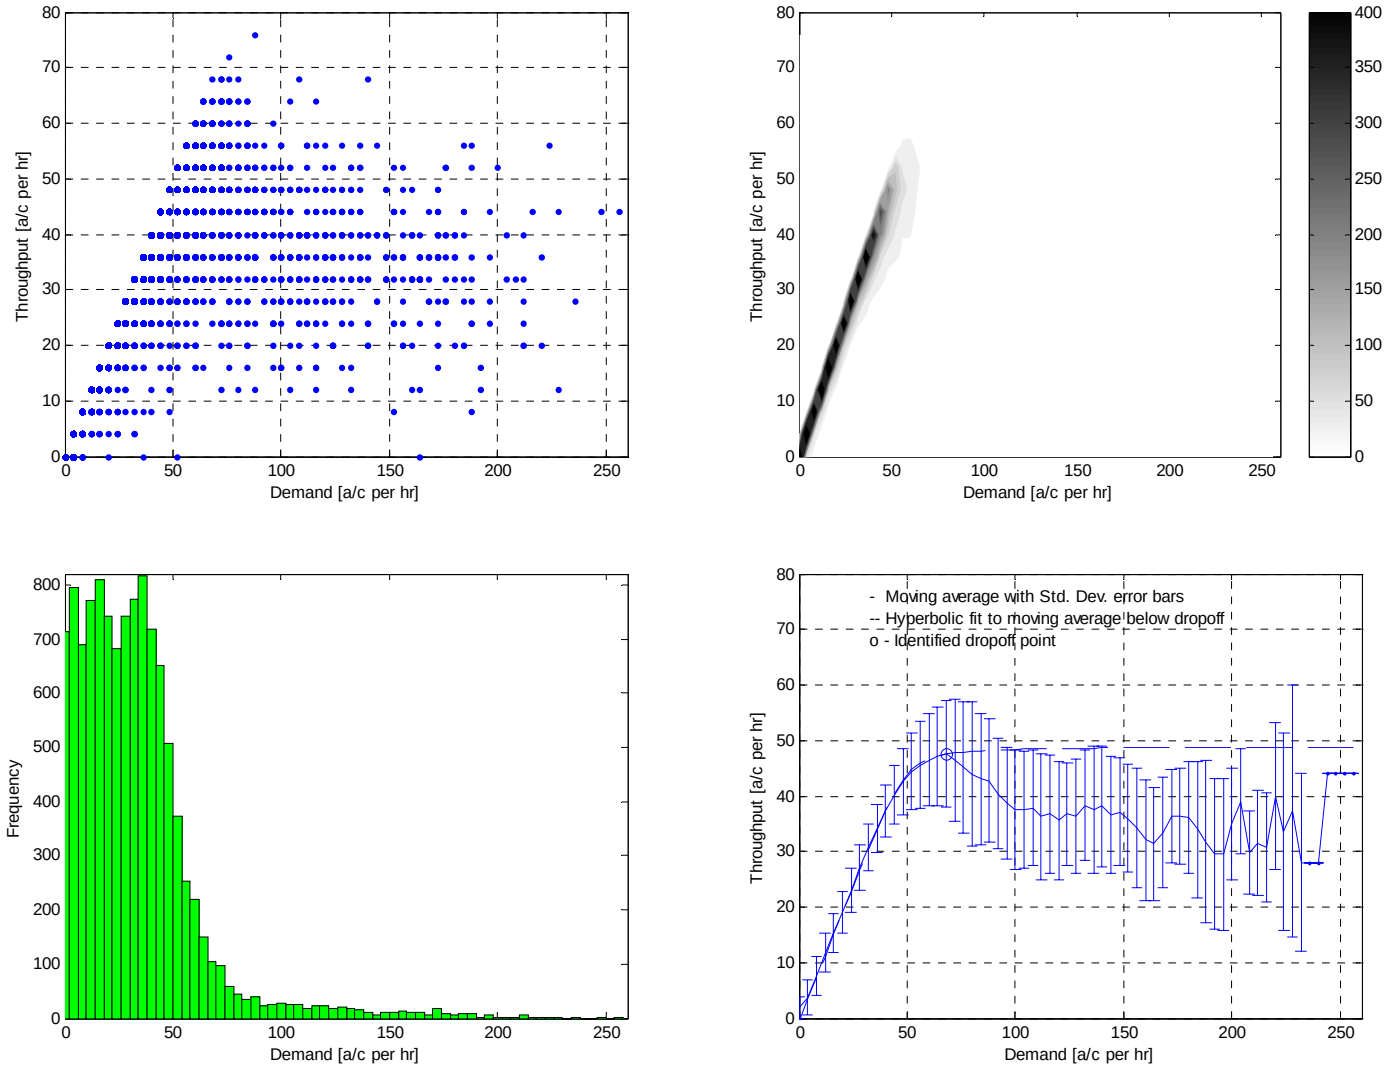

Figure B5. Airport throughput plotted agains7t demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	52	52
99	56	56
100	88	76
Hyperbolic Fit Asymptote		49
Hyperbolic Vertex 'a'		5
Moving Average Max	68	48
Moving Average Drop-off	68	48

B2. John F Kennedy International Airport JFK

AAR Frequency

The frequency of different AARs are plotted below. The data used is that from the summer of 2000, and from the winter, spring, and summer of 2001. The data used is per hour, as AARs are reported as hourly rates.

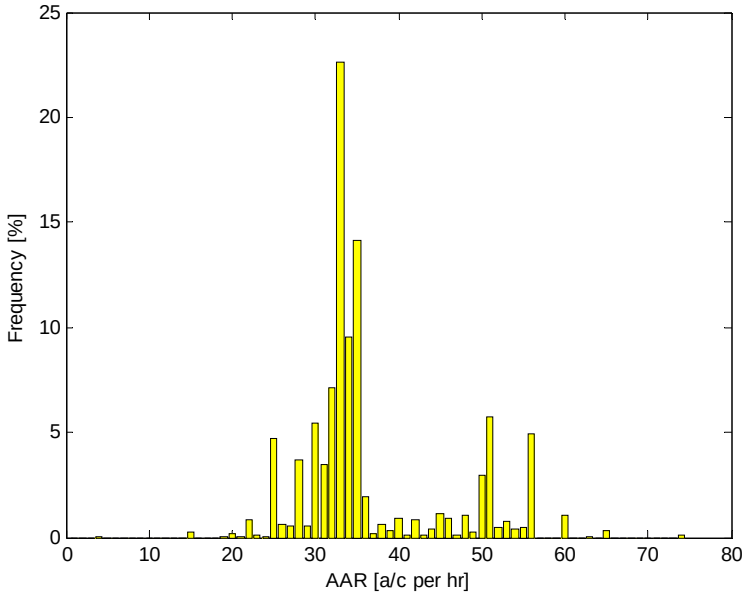

Figure B6. Frequency of reported AARs.

AARs of 33, 35 and 51 are examined in more detail below. These are the AARs reported on the days analyzed in the RTO 77 study.

B2.1. Airport Arrival Rate 33 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

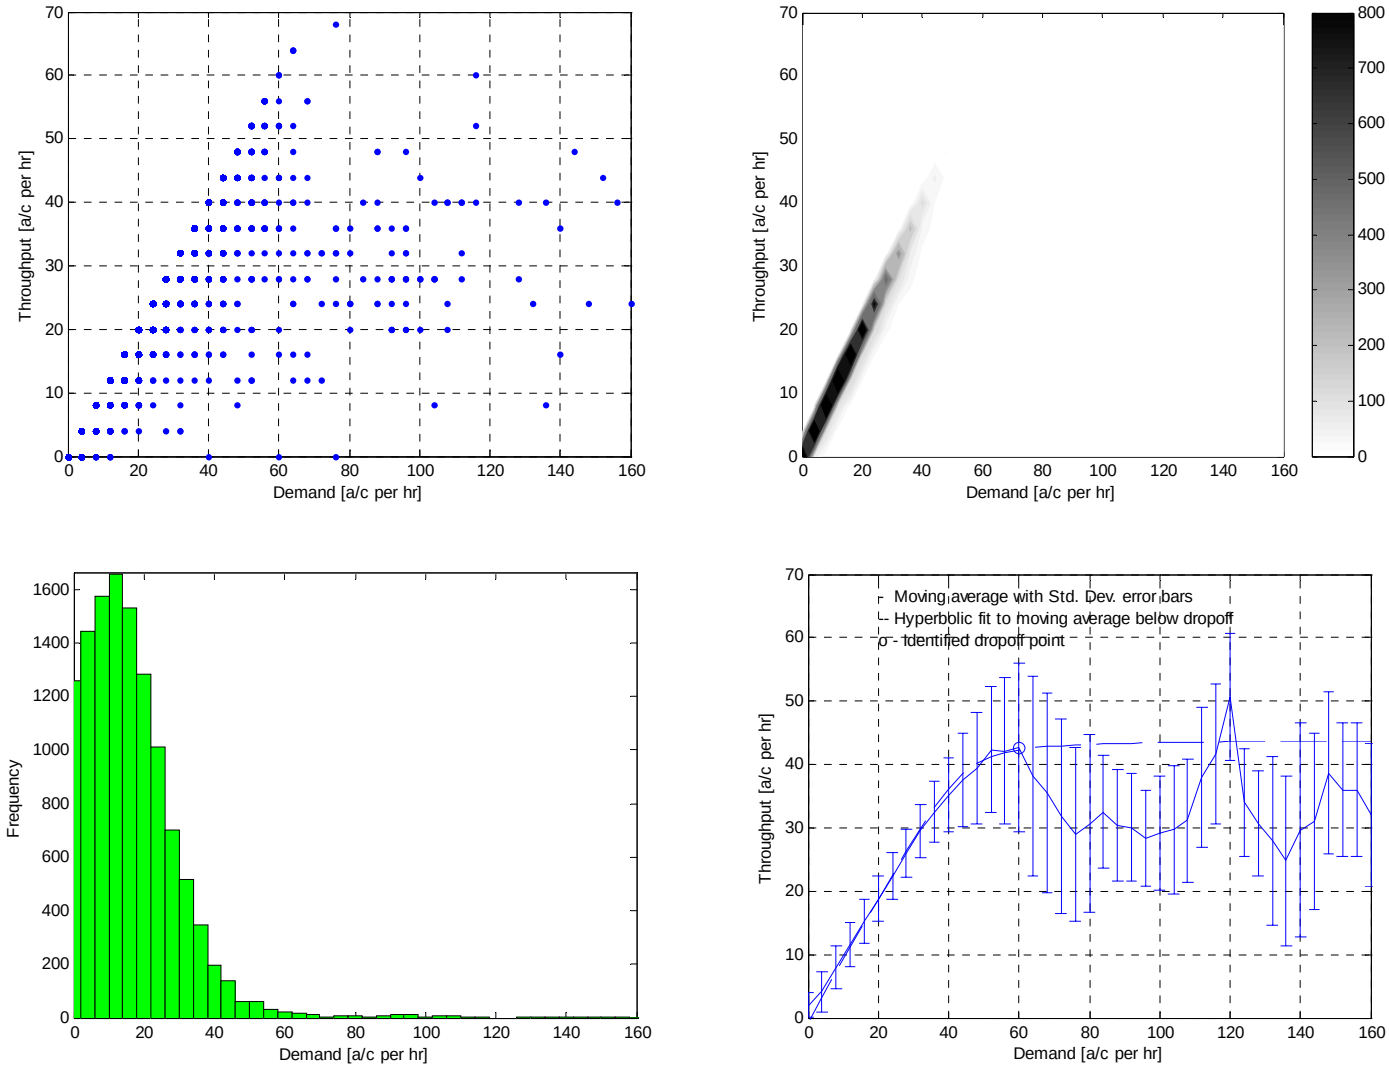

Figure B7. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	40	40
99	44	44
100	76	68
Exponential Fit Asymptote		44
Hyperbolic Vertex 'a'		5
Moving Average Max	60	43
Moving Average Drop-off	60	43

B2.2. Airport Arrival Rate 35 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

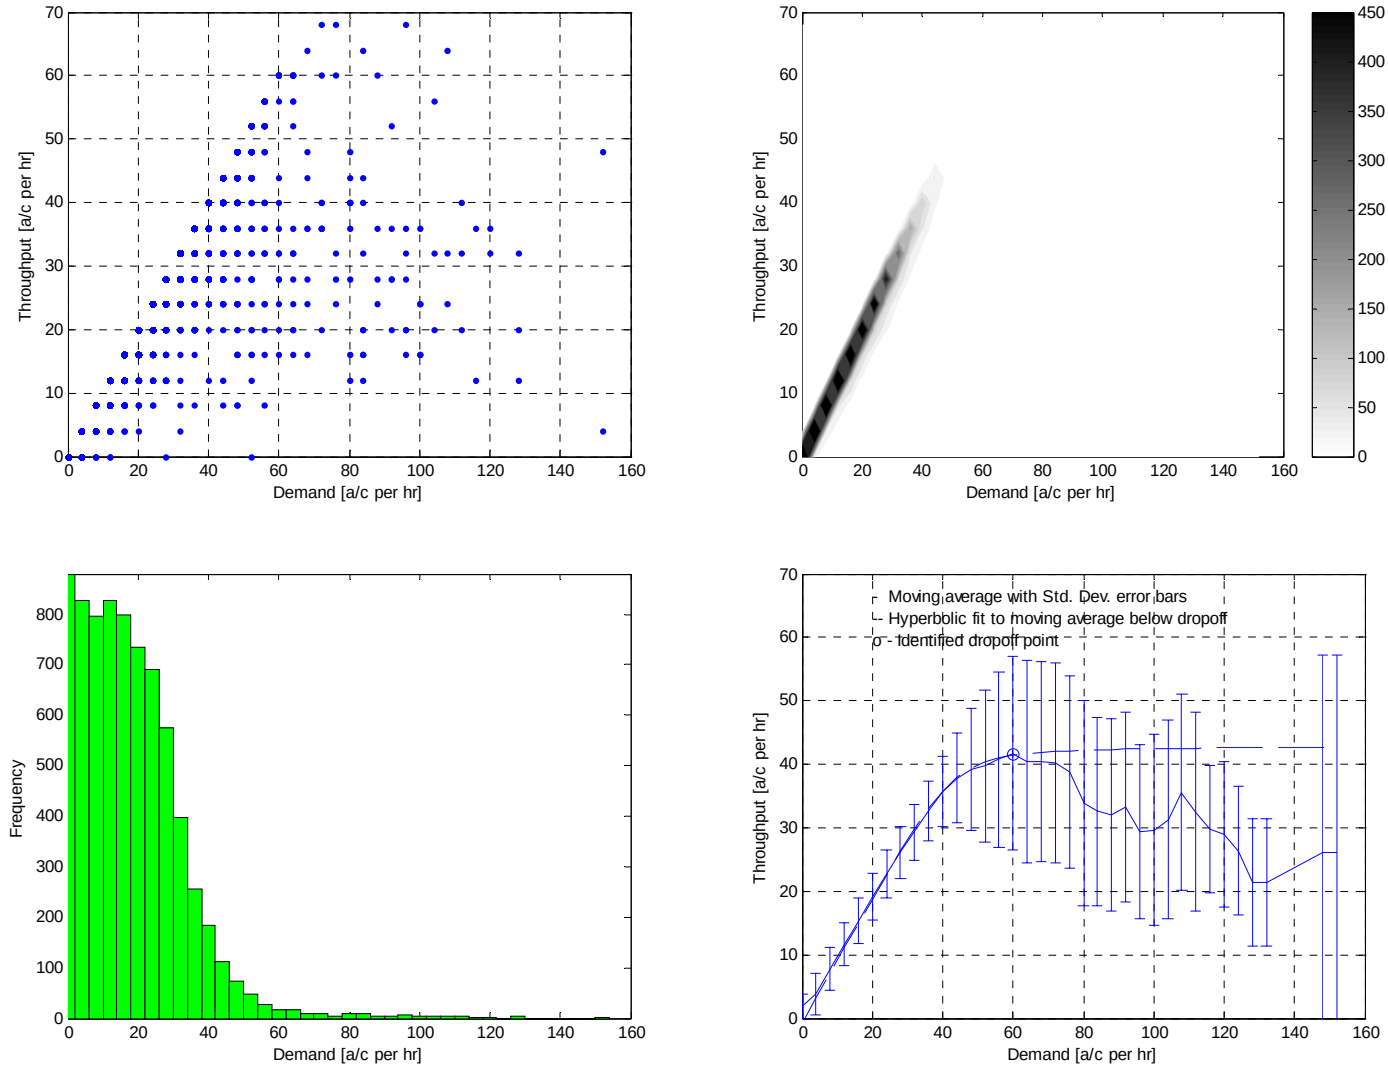

Figure B8. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	Demand [a/c per hr]	Throughput [a/c per hr]
98	44	44
99	48	48
100	72	68
Hyperbolic Fit Asymptote		43
Hyperbolic Vertex 'a'		5
Moving Average Max	60	42
Moving Average Drop-off	60	42

B2.3. Airport Arrival Rate 51 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

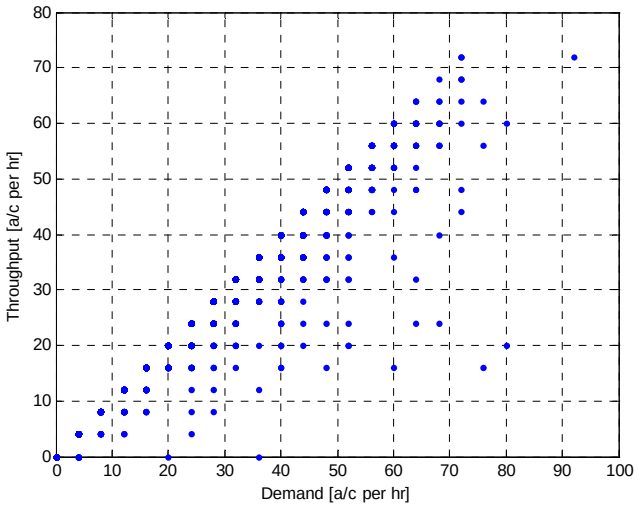

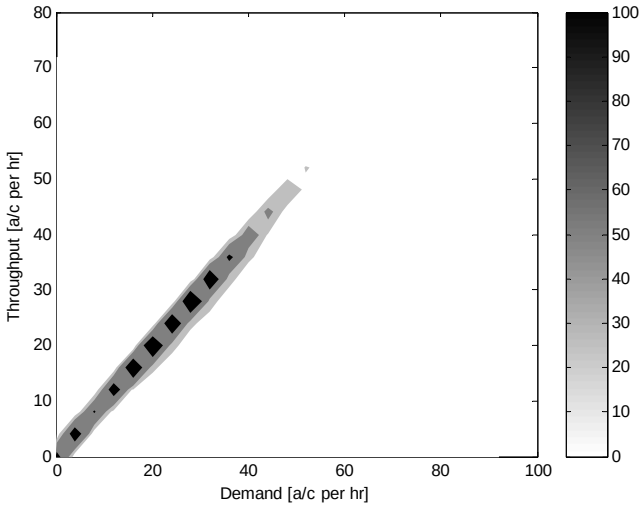

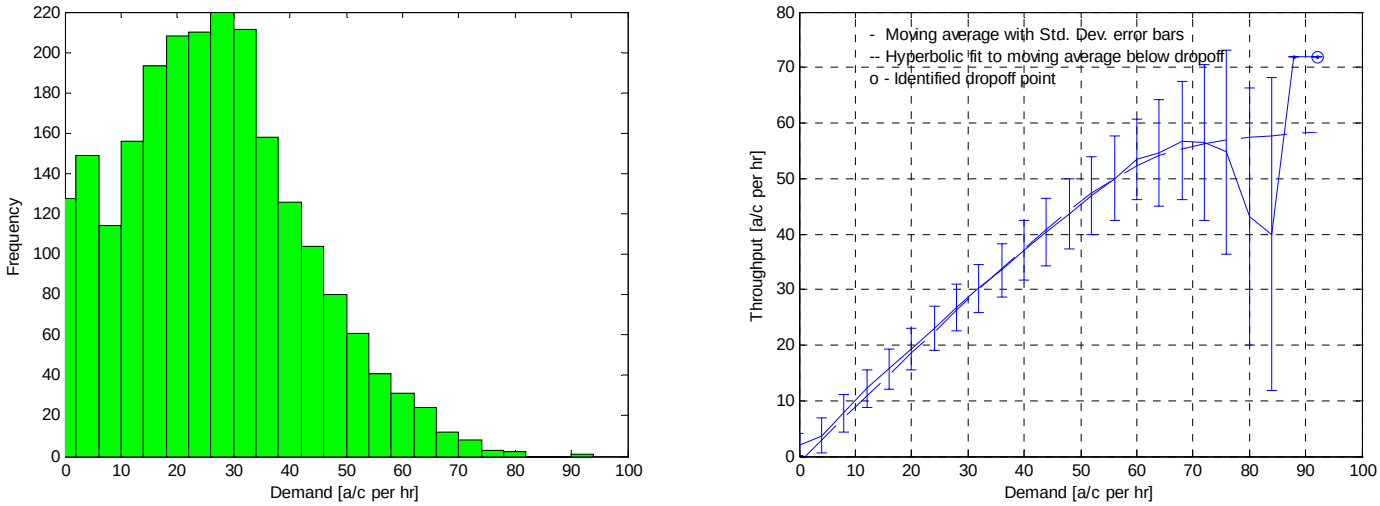

Figure B9. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	56	56
99	60	60
100	72	72
Exponential Fit Asymptote		60
Hyperbolic Vertex 'a'		7
Moving Average Max	88	72
Moving Average Drop-off	92	72

B3. LaGuardia Airport LGA

AAR Frequency

The frequency of different AARs are plotted below. The data used is that from the summer of 2000, and from the winter, spring, and summer of 2001. The data used is per hour, as AARs are reported as hourly rates.

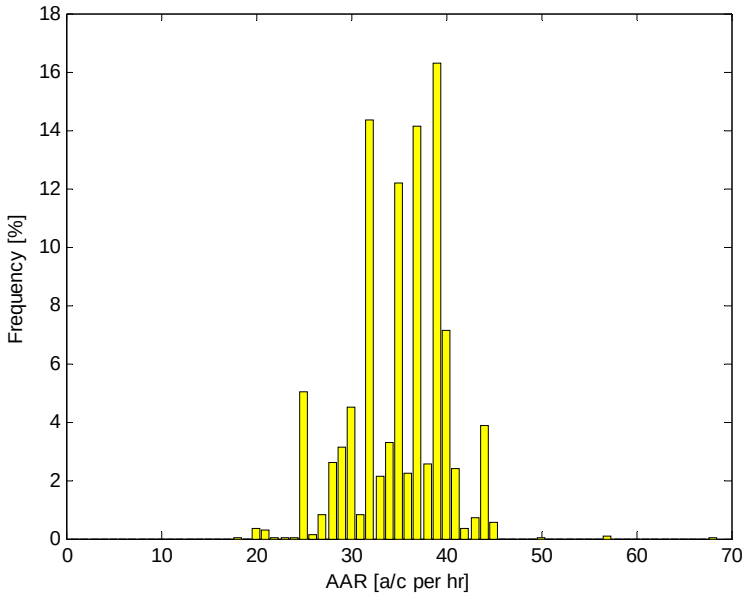

Figure B10. Frequency of reported AARs.

AARs of 31, 34, 39 and 42 are examined in more detail below. These are some of the AARs reported on the days analyzed in the RTO 77 study.

B3.1. Airport Arrival Rate 31 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

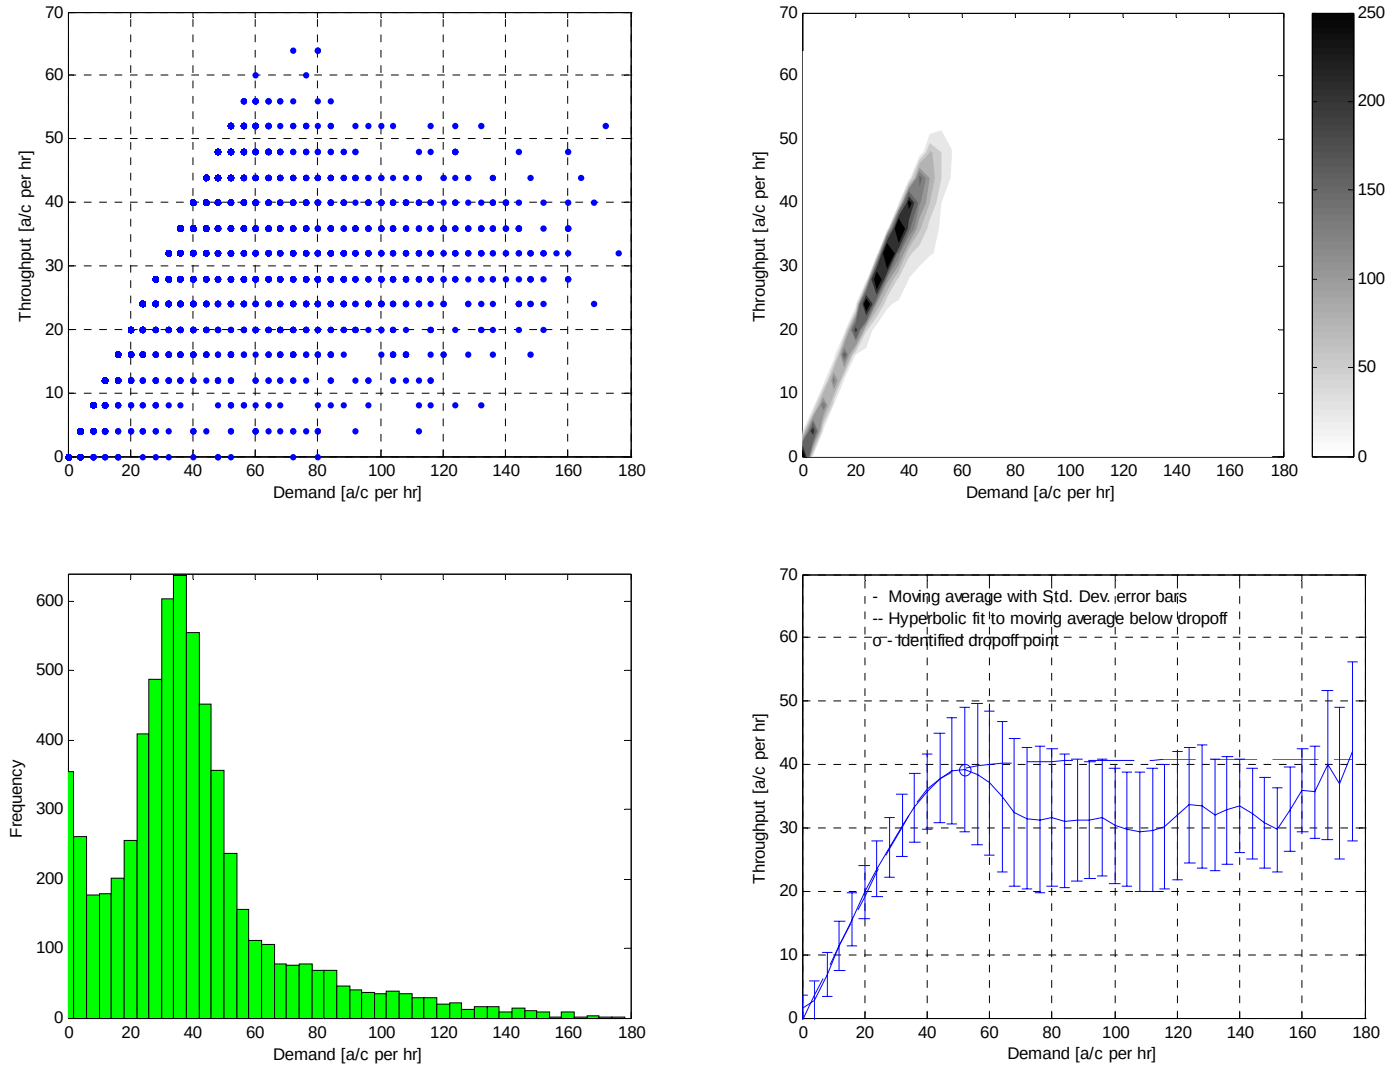

Figure B11. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	48	48
99	52	52
100	72	64
Exponential Fit Asymptote		41
Hyperbolic Vertex 'a'		4
Moving Average Max	52	39
Moving Average Drop-off	52	39

B3.2. Airport Arrival Rate 34 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

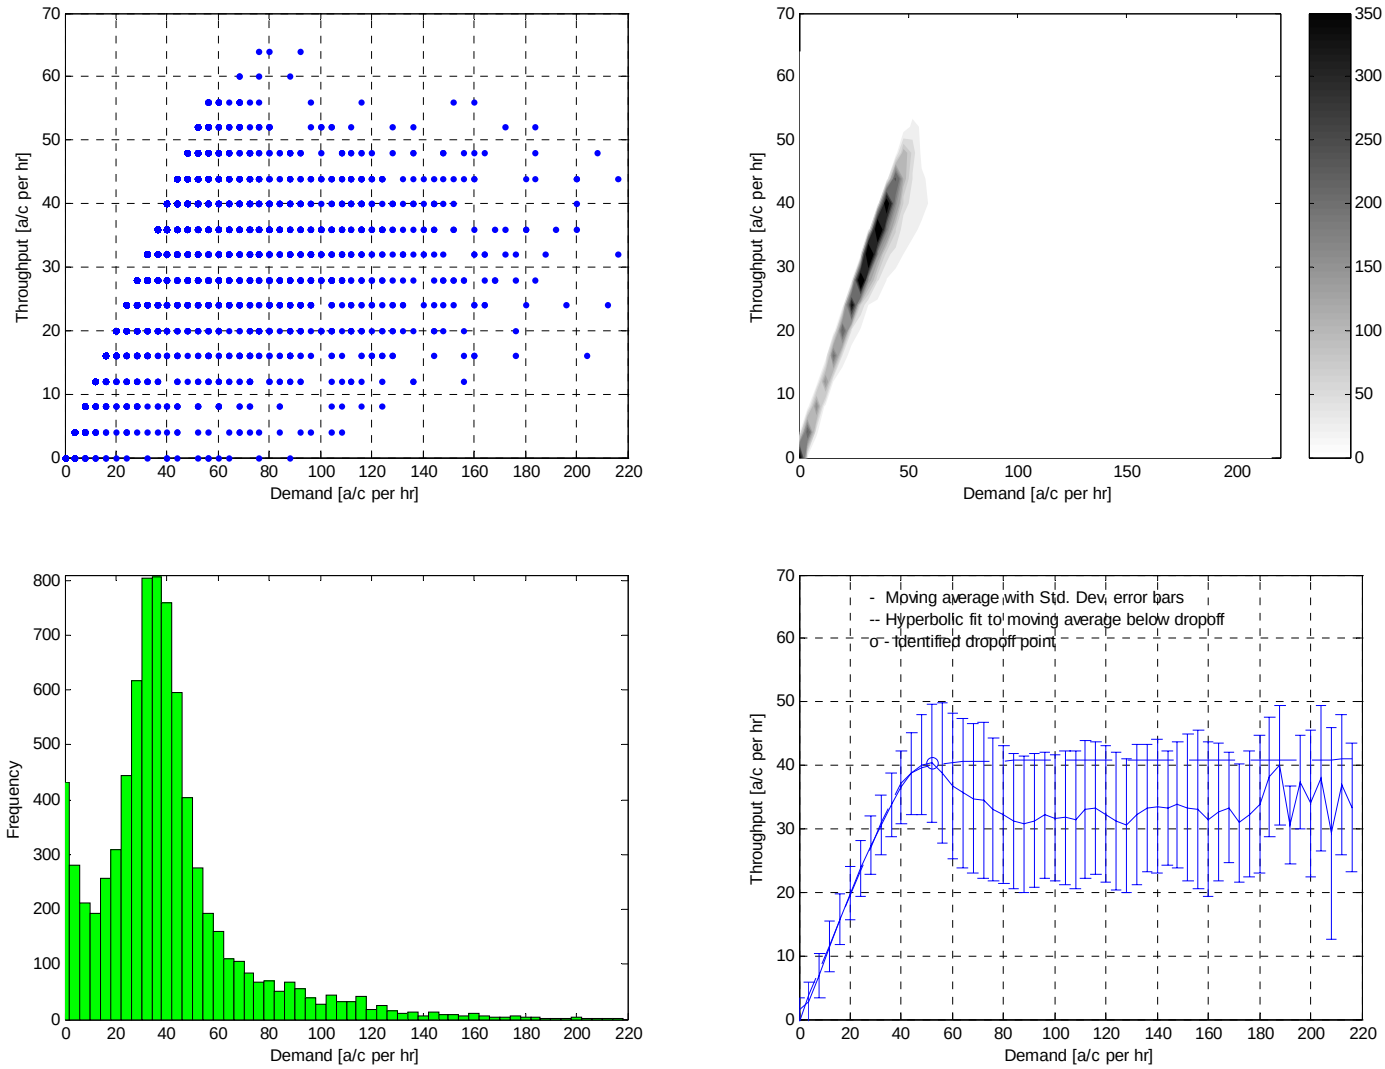

Figure B12. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	Demand [a/c per hr]	Throughput [a/c per hr]
98	48	48
99	52	52
100	76	64
Exponential Fit Asymptote		41
Hyperbolic Vertex 'a'		3
Moving Average Max	52	40
Moving Average Drop-off	52	40

B3.3. Airport Arrival Rate 39 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

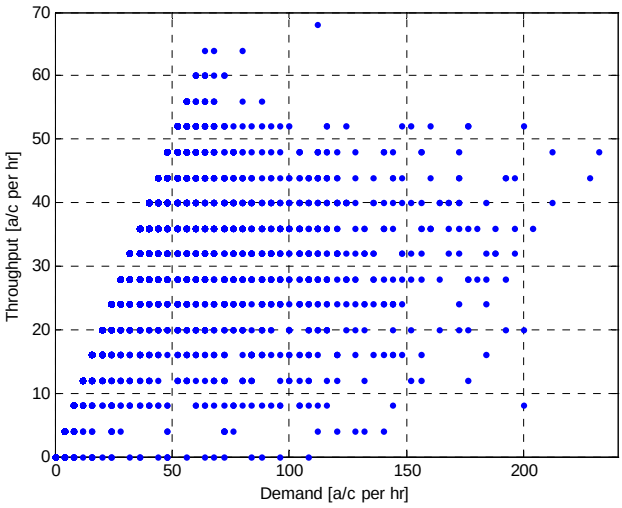

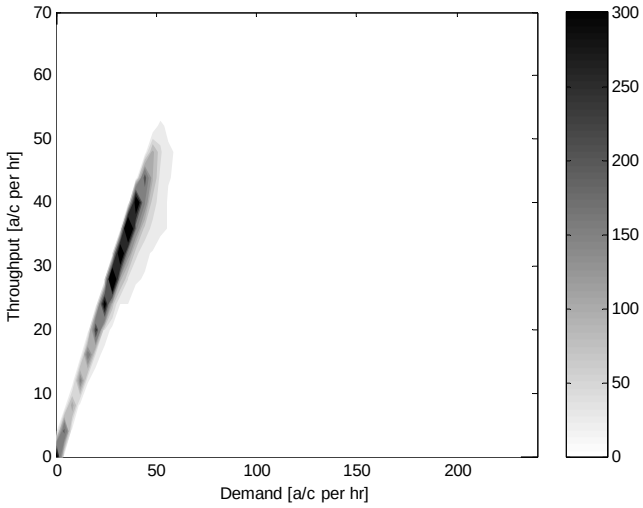

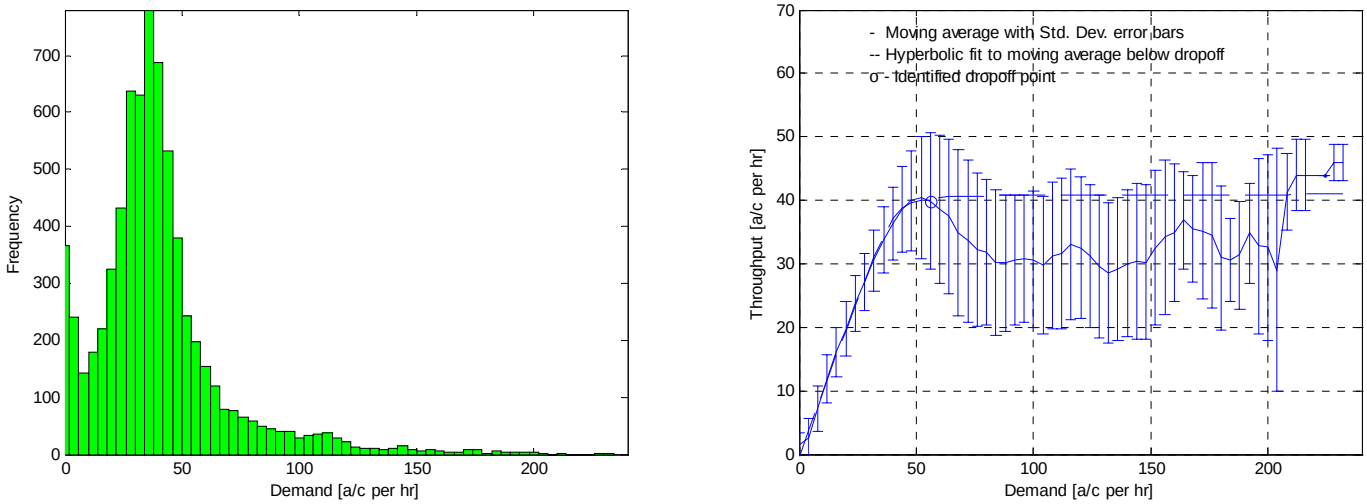

Figure B13. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	52	52
99	52	52
100	112	68
Exponential Fit Asymptote		41
Hyperbolic Vertex 'a'		3
Moving Average Max	52	40
Moving Average Drop-off	56	39

B3.4. Airport Arrival Rate 42 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

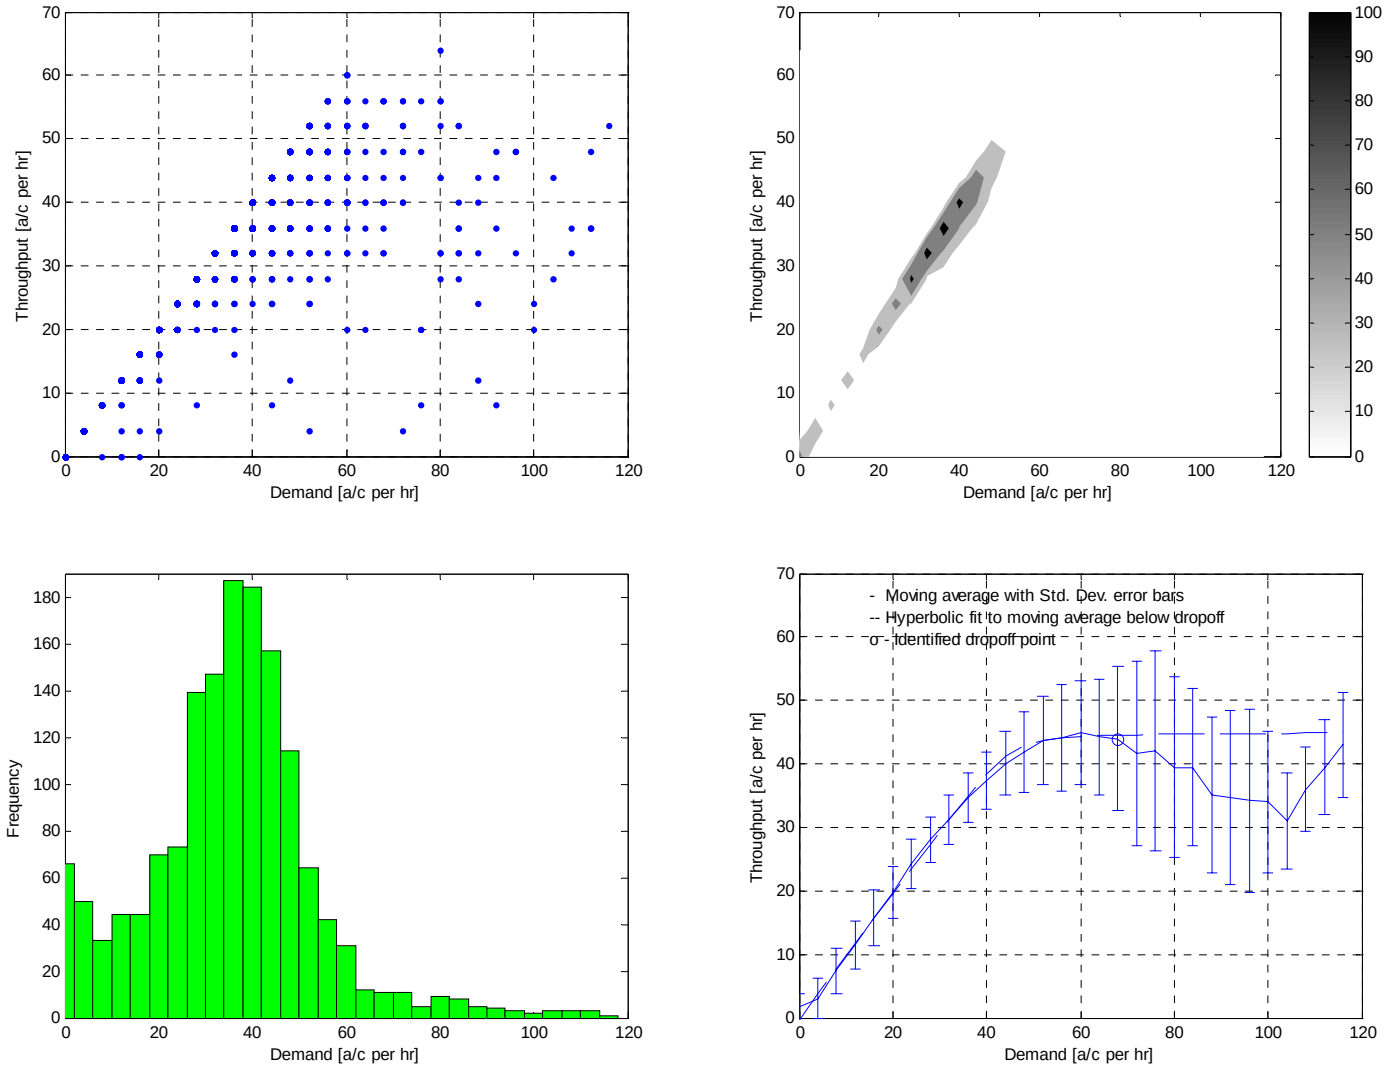

Figure B14. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	52	52
99	56	56
100	80	64
Exponential Fit Asymptote		45
Hyperbolic Vertex 'a'		3
Moving Average Max	60	45
Moving Average Drop-off	68	44

B4. Philadelphia International Airport PHL

AAR Frequency

The frequency of different AARs are plotted below. The data used is that from the summer of 2000, and from the winter, spring, and summer of 2001. The data used is per hour, as AARs are reported as hourly rates.

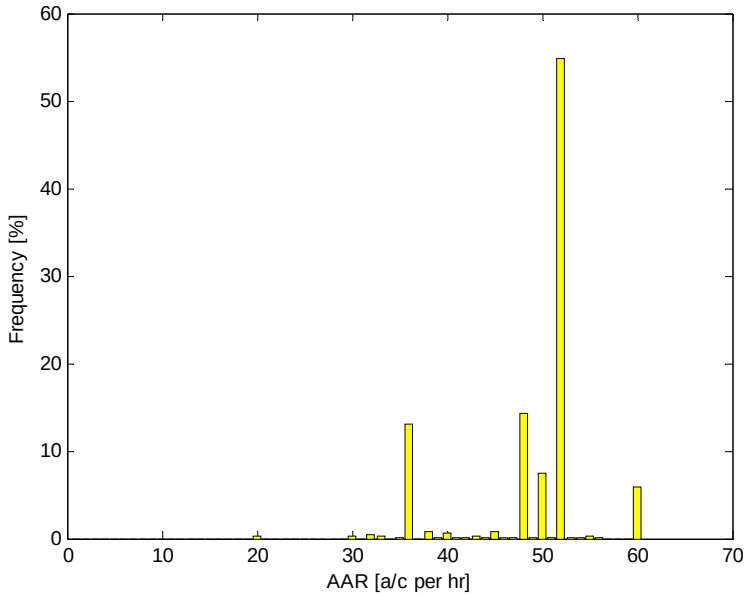

Figure B15. Frequency of reported AARs for this configuration.

AARs of 36 and 52 are examined in more detail below. These are the AARs reported on the days analyzed in the RTO 77 study.

B4.1. Airport Arrival Rate 36 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

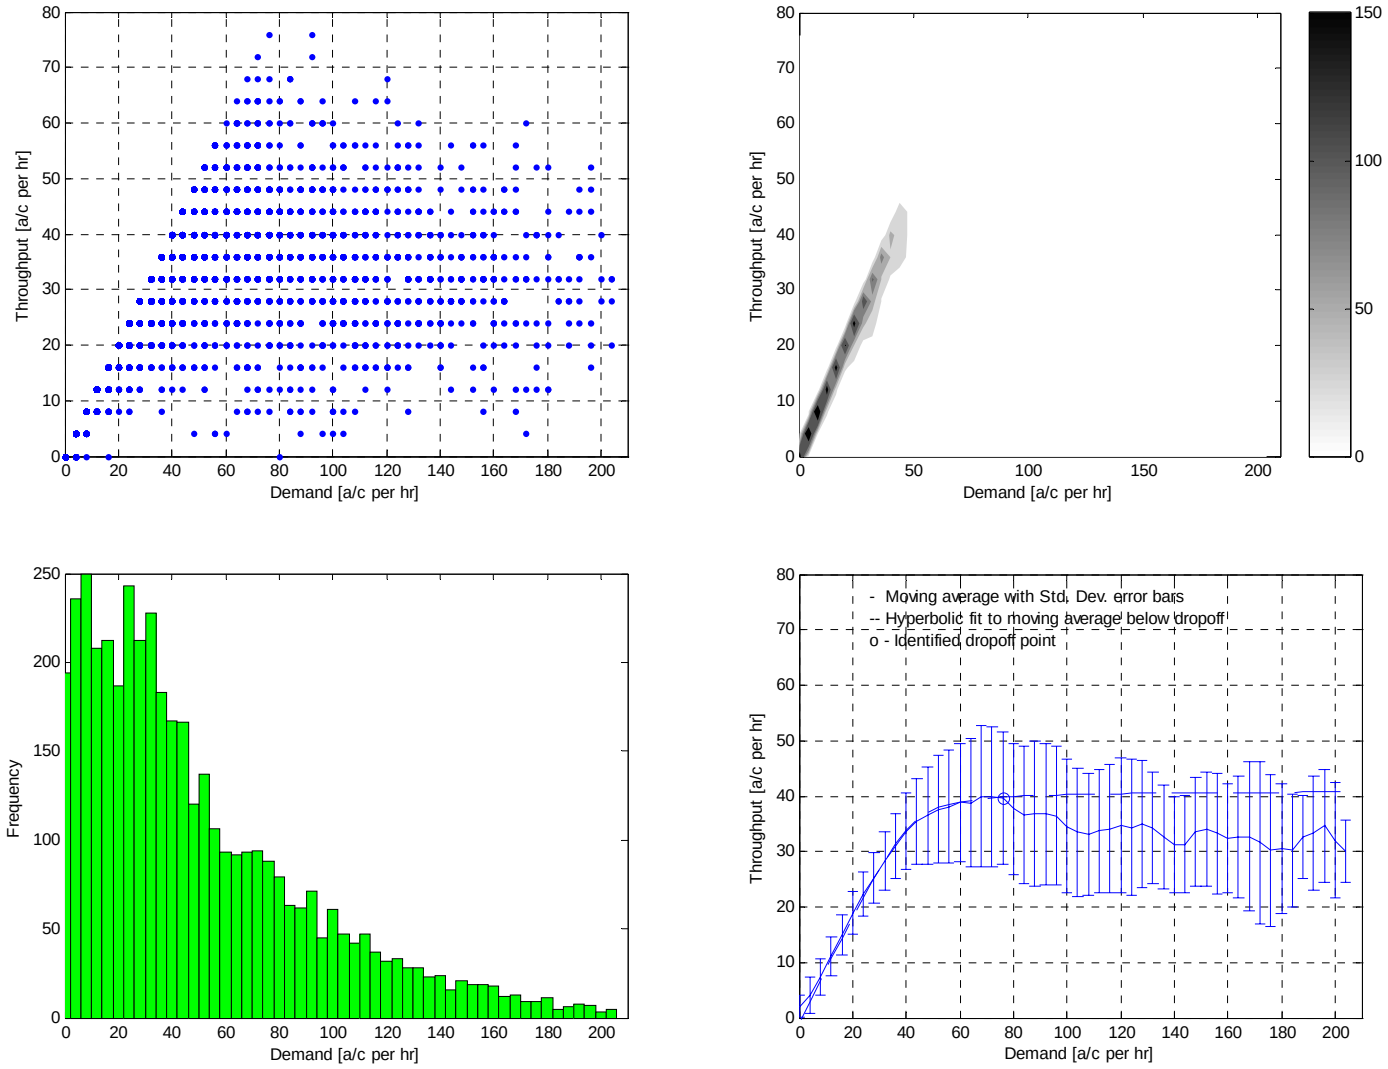

Figure B16. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	56	56
99	60	60
100	76	76
Exponential Fit Asymptote		41
Hyperbolic Vertex 'a'		6
Moving Average Max	72	40
Moving Average Drop-off	76	40

B4.2. Airport Arrival Rate 52 aircraft per hour

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

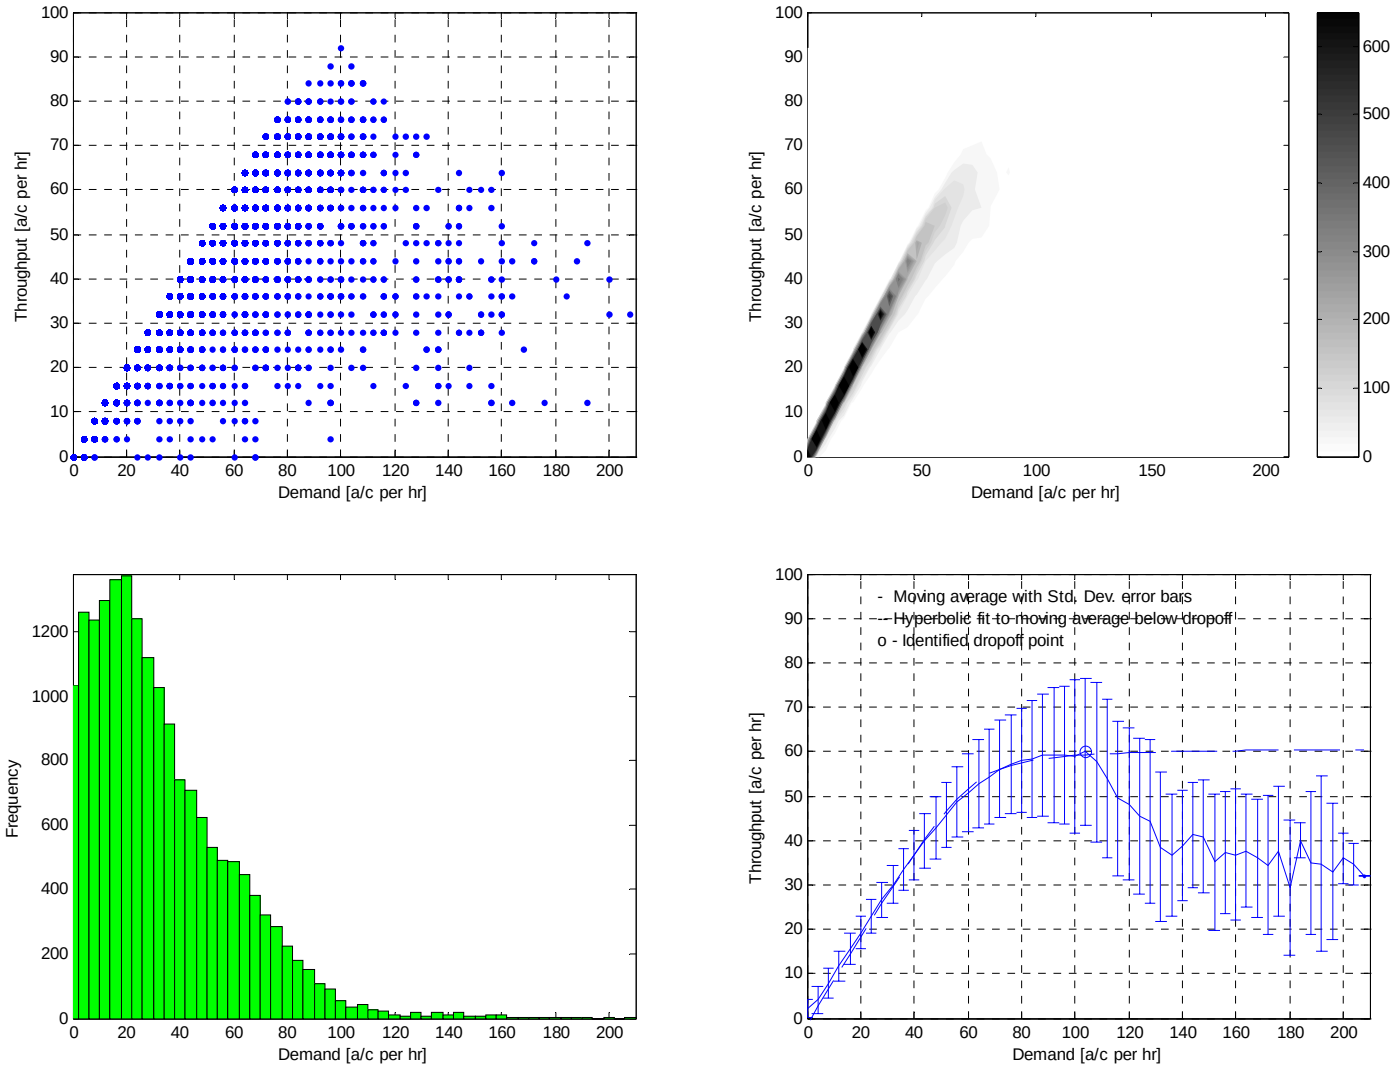

Figure B17. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	68	68
99	72	72
100	100	92
Exponential Fit Asymptote		61
Hyperbolic Vertex 'a'		8
Moving Average Max	104	60
Moving Average Drop-off	104	60

B5. Teterboro Airport TEB

AAR Frequency

The frequency of different AARs are plotted below. The data used is that from the summer of 2000, and from the winter, spring, and summer of 2001. The data used is per hour, as AARs are reported as hourly rates.

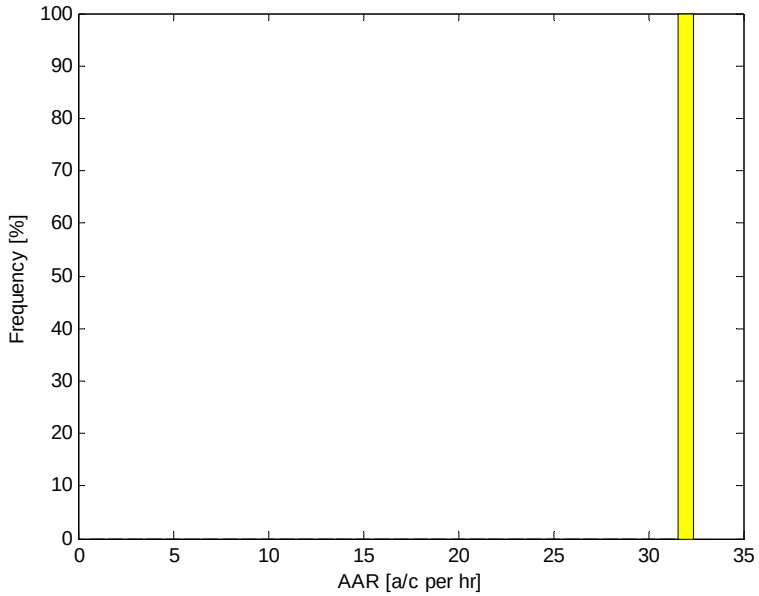

Figure B18. Frequency of reported AARs.

Throughput versus Demand

In order to identify airport capacity, throughput is plotted against arrival demand in the charts below.

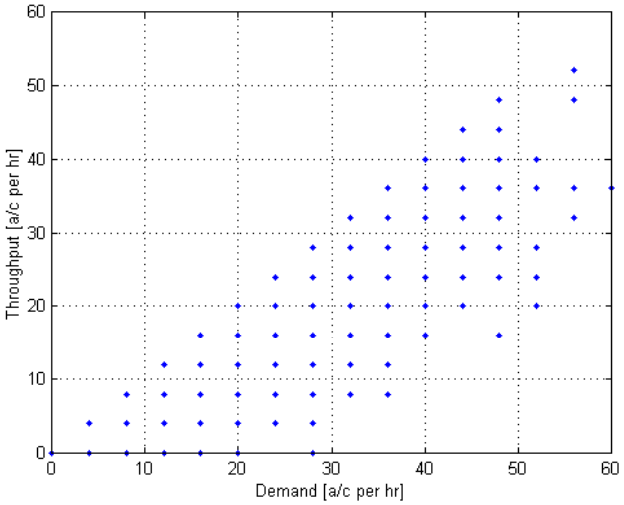

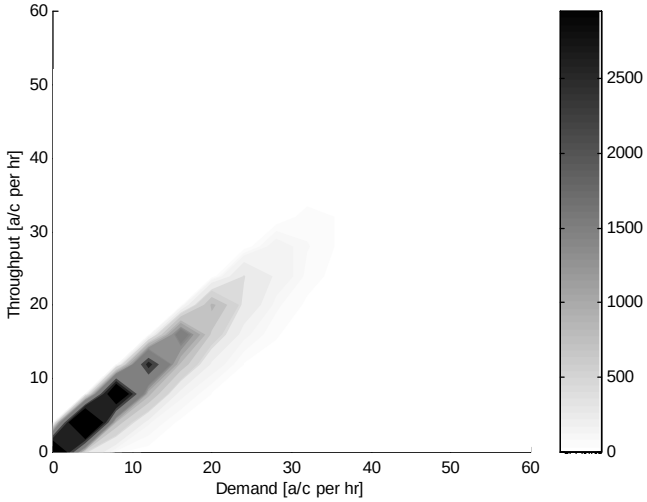

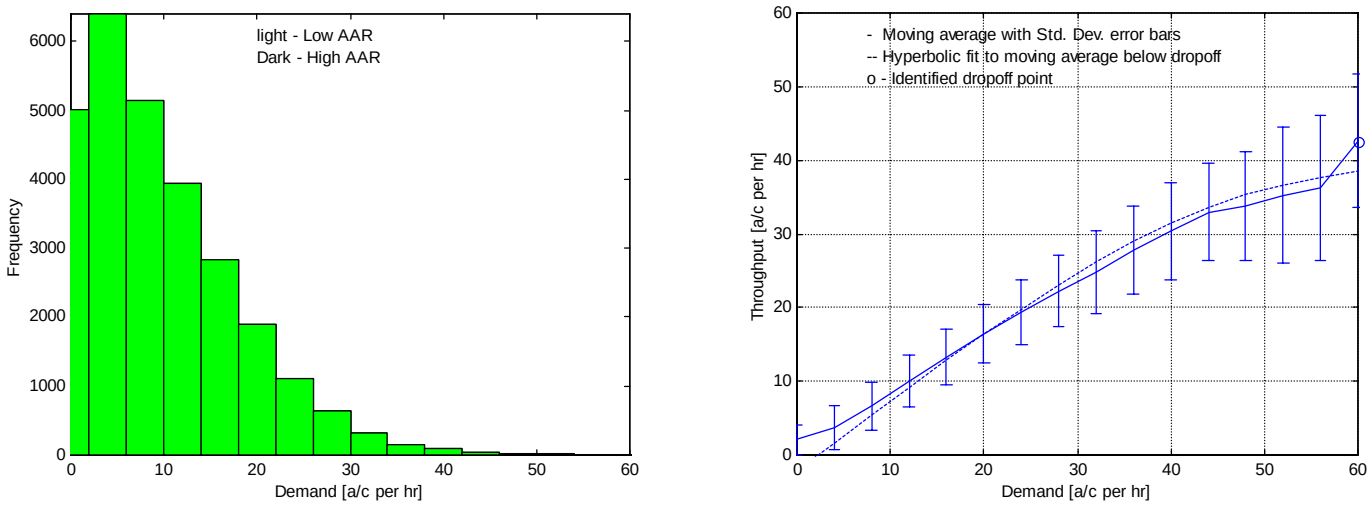

Figure B19. Airport throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	28	28
99	28	28
100	56	52
Exponential Fit Asymptote		43
Hyperbolic Vertex 'a'		9
Moving Average Max	60	43
Moving Average Drop-off	No drop-off	

B6. New York TRACON N90

The data from EWR, JFK, LGA and TEB is combined for those days for which all data sets have data. In the first set of results all the data is plotted. In subsequent sections the results are presented for specific combinations of airport configurations.

AAR Frequency

The frequency of different AARs are plotted below. Hourly data is used because AARs are generally reported as hourly rates.

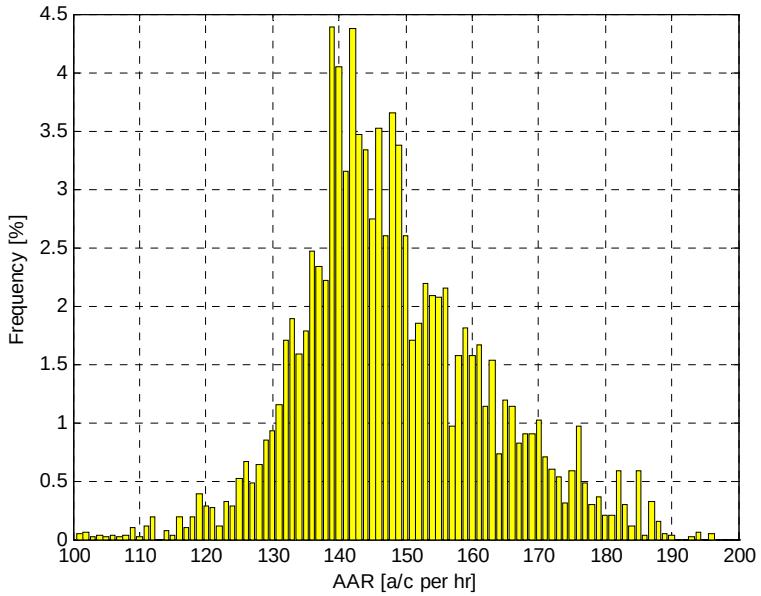

Figure B20. Frequency of reported AARs.

B6.1. AAR of 142 aircraft per hour

Throughput versus Demand

In order to identify TRACON capacity, throughput is plotted against arrival demand in the charts below.

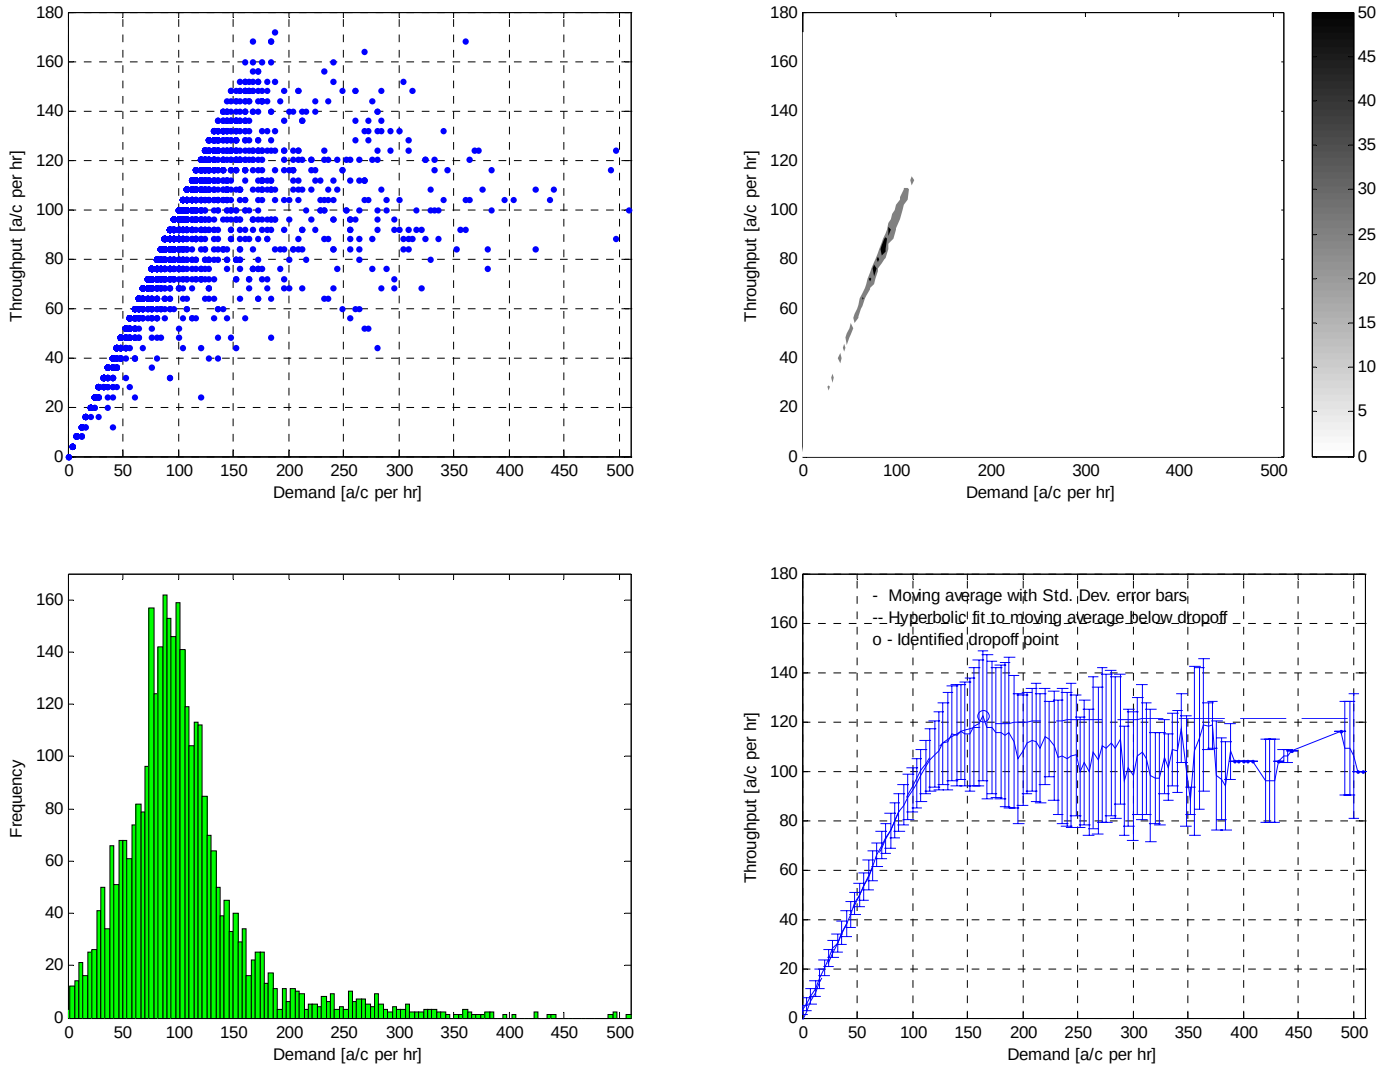

Figure B21. TRACON throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	144	144
99	148	148
100	188	172
Exponential Fit Asymptote		122
Hyperbolic Vertex 'a'		12
Moving Average Max	164	122
Moving Average Drop-off	164	122

B6.2. AAR of between 139 and 142 aircraft per hour

Throughput versus Demand

In order to identify TRACON capacity, throughput is plotted against arrival demand in the charts below. A frequency contour plot is not included as the frequency of data points in the 4 aircraft per hour by 4 aircraft per hour bins is never greater than 1.

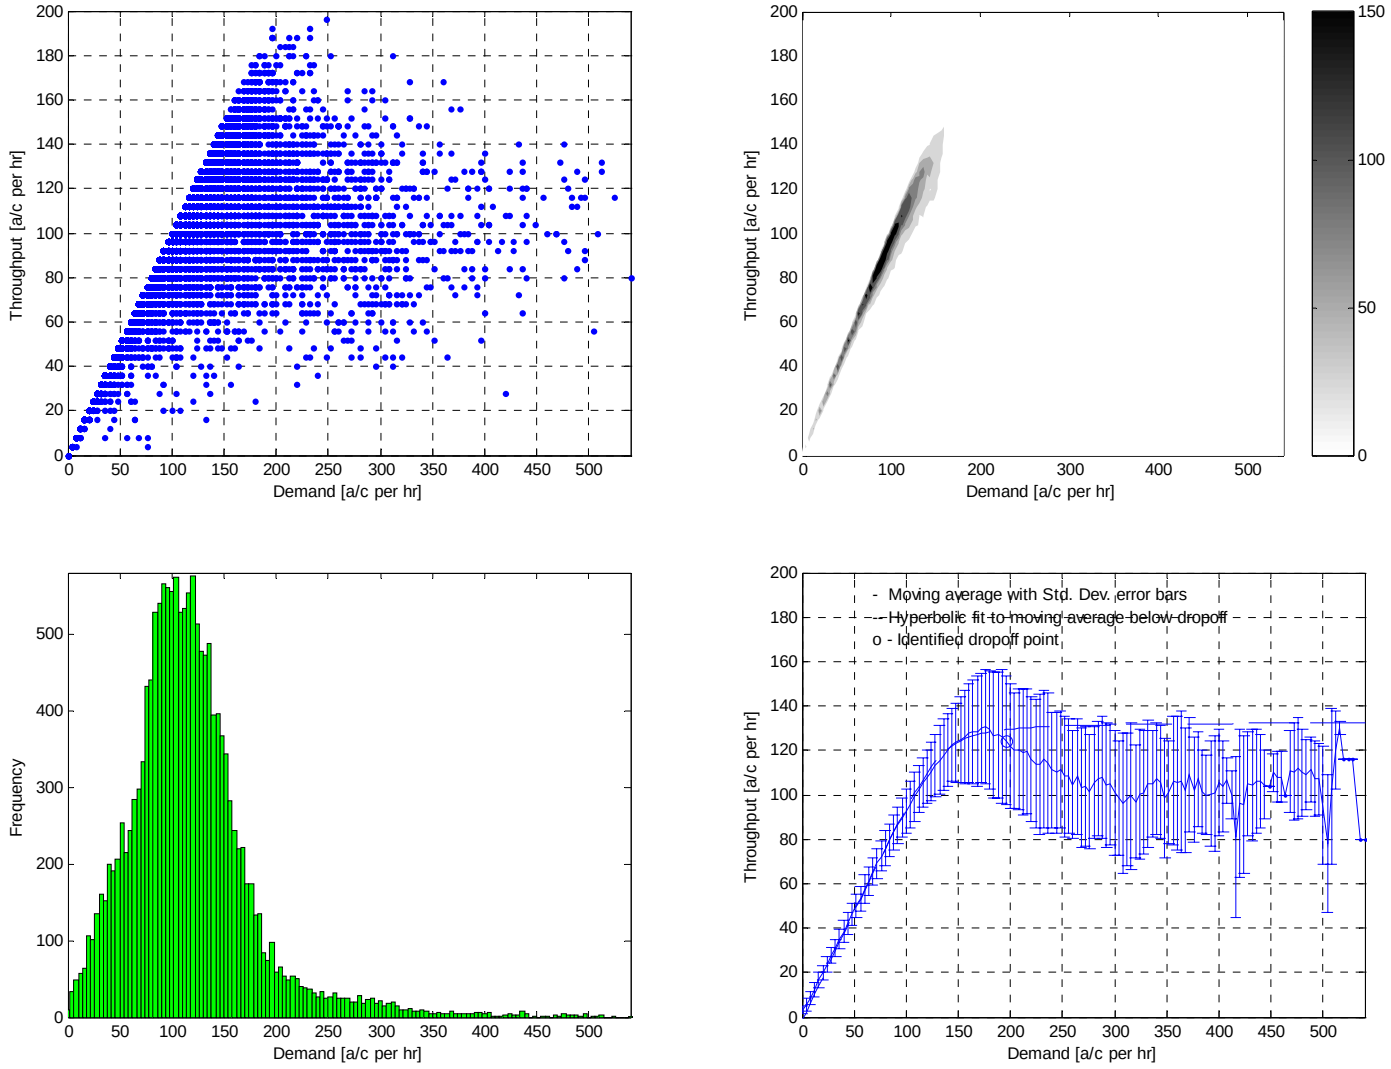

Figure B22. TRACON throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	Demand [a/c per hr]	Throughput [a/c per hr]
98	152	152
99	160	160
100	248	196
Exponential Fit Asymptote		133
Hyperbolic Vertex 'a'		14
Moving Average Max	176	131
Moving Average Drop-off	196	124

B6.3. EWR AAR of 44 ac/hr, JFK AAR of 33 ac/hr, LGA AAR of 39 ac/hr, and TEB AAR of ac/hr (giving N90 AAR of 148 ac/hr).

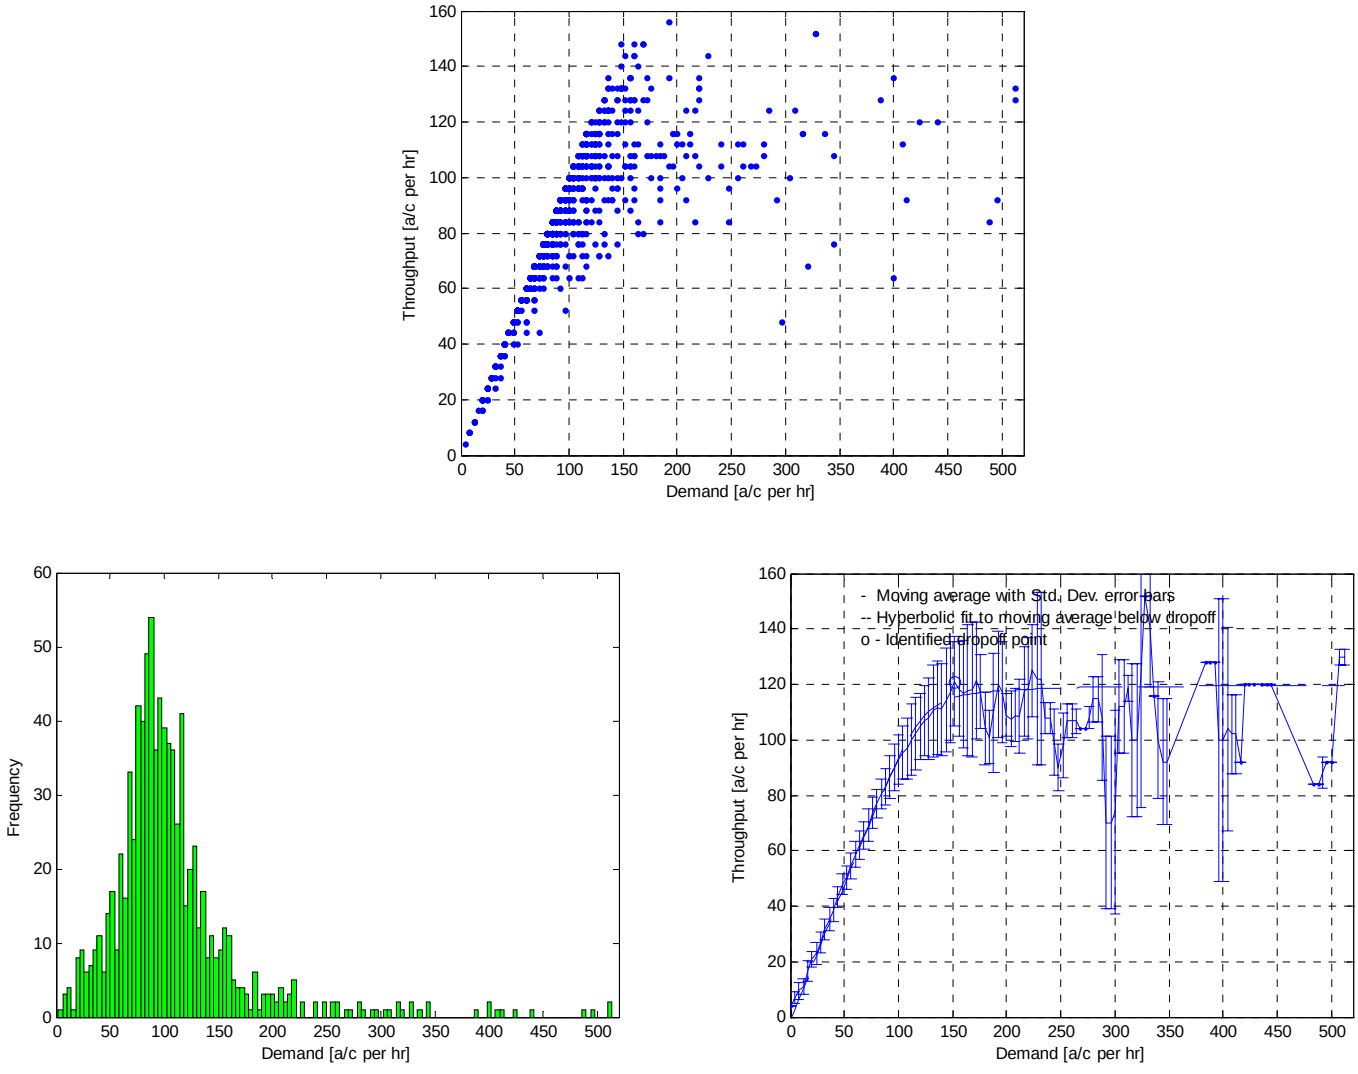

Figure B23. TRACON throughput plotted against demand – raw data plot, contour frequency plot, demand frequency plot, and moving average plot with hyperbolic curve fit.

Percentile of Throughput	**Demand** [a/c per hr]	**Throughput** [a/c per hr]
98	136	136
99	152	144
100	192	156
Exponential Fit Asymptote		120
Hyperbolic Vertex 'a'		12
Moving Average Max	152	121
Moving Average Drop-off	152	121

B7. Observations

A number of observations were made from the plots of throughput versus demand. These are as follows:

1. The majority of the data points lie along the leading edge of the curve (45 º radial). Generally there are more points on the leading edge at lower demand than higher. This suggests that the points of very higher demand occur relatively less frequently. It also suggests that data points in which demand is high and throughput is low are infrequent. Thus, demand is generally matched by capacity.
2. All the curves show a maximum throughput. However, the degree to which the curves appear to saturate to this maximum varies. It is suggested that the degree to which a configuration saturates indicates the degree to which it is constrained by runways. This is supported by a number of observations, as follows:
    2.1. Both high and low capacity (AAR) configurations at PHL show clear saturation to the maximum throughput. According to NASA McTMA researchers, PHL is highly constrained by runway constraints.
    2.2. The high capacity configuration at JFK shows no saturation. At high capacity JFK is expected to have few runway constraints.
    2.3. EWR and LGA show more saturation than JFK, but not as much as PHL. Both airports are constrained by runways to a certain extent, but are also expected to be constrained by the airspace.
    2.4. The low capacity configurations (at JFK, LGA, and PHL) show more saturation than the higher capacity configurations. Runway limitations are likely to be a greater constraint in low capacity configurations than high capacity configurations, supporting the suggestion that the degree of saturation relates to runway constraints.
3. Almost all curves show a drop-off in throughput at high demand. The location of this drop-off relative to the peak in throughput varies, however. It is suggested that the position of the drop-off relative to the peak in throughput is an indication of the ability to maintain pressure on the runway, when at maximum capacity, through such techniques as a managed reservoir. This is supported by a number of observations, as follows:
    3.1. The drop-off at PHL is later than at the N90 airports. PHL does operate a managed reservoir, according to NASA McTMA researchers, and is thus expected to be able to maintain throughput at capacity for longer.
    3.2. The drop-off in capacity occurs directly after the peak in throughput at all airports within N90. N90 does not operate a managed reservoir, and is unable to hold aircraft near their destination airports because of airspace constraints. Instead the flow must be held further upstream, effectively switching off the flow at the airport. This is expected to lead to a sudden drop off immediately after the peak in throughput, as observed.

4. The degree to which the throughput drops off also varies by configuration. It is suggested that this indicates how efficiently aircraft can be held for an airport. This is similar to the managed reservoir, but applies to operations further upstream as well. The suggestion is supported by the following observations:

    4.1. The level to which the throughput drops remains approximately the same between the high and low capacity configurations (at JFK, LGA and PHL). This suggests that these levels are a function of constraints upstream from the airport, such as holding capacity upstream.

    4.2. The drop-off at JFK, LGA and EWR is severe (most severe at JFK). N90 is not able to hold aircraft near the airport, and consequently, when demand exceeds capacity, the throughput is expected to drop significantly because aircraft must be held well outside N90.

    4.3. The drop-off in throughput for all airports within N90 combined is also significant, although not as severe as that for each airport respectively.

Appendix C: ASPM Capacity Utilization Analysis

Following is a capacity utilization analysis of the four primary airports in New York (EWR, JFK, LGA and TEB), and of Philadelphia airport (PHL). As in Appendix B, the analysis is performed using the Aviation System Performance Metrics (ASPM) database, obtained through NASA and the FAA (http://www.apo.data.faa.gov/faamatsall.HTM). The results of the capacity utilization analysis are separated by airport. In order to accurately analyze the capacity utilization of the airports the data analyzed must include the periods of highest demand. Consequently data from January to August of 2001 was chosen for processing. The summer months are the period of highest demand, so data from the summer of 2000 was also included to enhance the data from the summer of 2001. This choice of data ensures that the maximum airport capacity utilizations determined in the analysis are representative of the actual maximum capacity utilization of the airports.

C1. Newark International Airport EWR

Capacity utilization is plotted for a single airport configuration below, filtering for high demand periods only, followed by a plot of capacity utilization for all airport configurations combined, with no filtering for high demand. The configuration plotted is one of the commonly operated configurations at EWR.

C1.1. Configuration 3 – 22L | 22R

Actual arrival throughput is plotted against specified AAR in order to identify the degree to which the available capacity is being utilized. Only data with demand high enough to utilize the airport's maximum capacity is plotted. This thus excludes data points from the chart for which throughput was low simply because demand was low. Demand is identified as being high enough to utilize the airport's maximum capacity according to a specified demand threshold. This demand threshold is specified as the lowest demand for which the throughput reaches the 98^{th} percentile of throughput identified in the capacity analysis in Appendix B. In the case below this is 56 aircraft per hour.

The frequency of the data is plotted alongside the raw data. Again, the frequency is calculated in 4 aircraft per hour by 4 aircraft per hour bins. The average of the data, for each 4 aircraft per hour bin of AAR, is also presented on the graph, with error bars of one standard deviation. In many cases the error is so small as not to be visible on the plot. A linear least squares fit to these averages is also included as a dotted line. The diagonal straight line from corner to corner represents 100% utilization of the AAR.

The charts below show quarter hourly data, although the values for throughput and AAR are multiplied by four in order to represent rates by the hour.

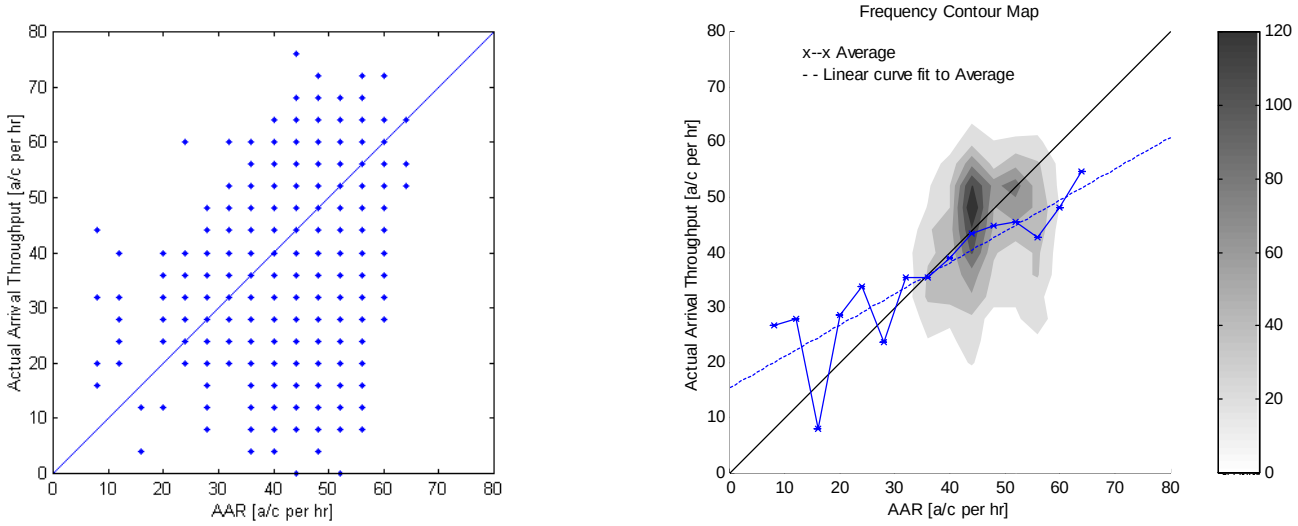

Figure C1. Airport utilization – raw data plot, and contour frequency plot.

The arrival rate above which capacity is underutilized by high demand is estimated as the point at which the linear fit to the averages crosses the diagonal 100% AAR utilization line. This is as follows:

Underutilization AAR	36 a/c per hr

C1.2. All EWR Configurations

The following plot shows all data points, with no filtering for high demand only:

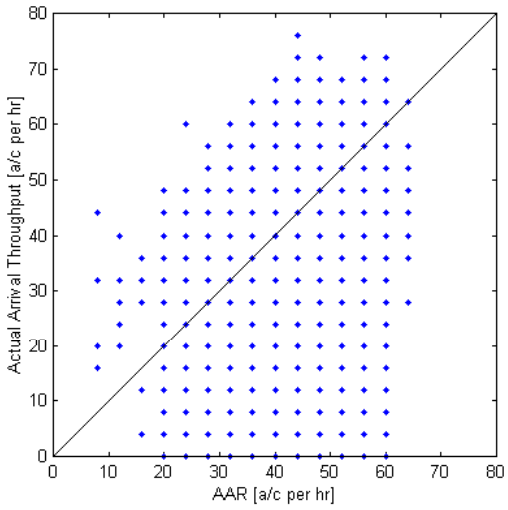

Figure C2. Airport utilization – raw data plot.

C2. John F Kennedy International Airport JFK

C2.1. Configuration 2 – 13L | 13R

In the following chart actual arrival throughput is plotted against specified AAR in order to identify the degree to which the available capacity is being utilized. Only data with demand high enough to utilize the airport's maximum capacity is plotted. The threshold according to which demand is filtered is 44 aircraft per hour.

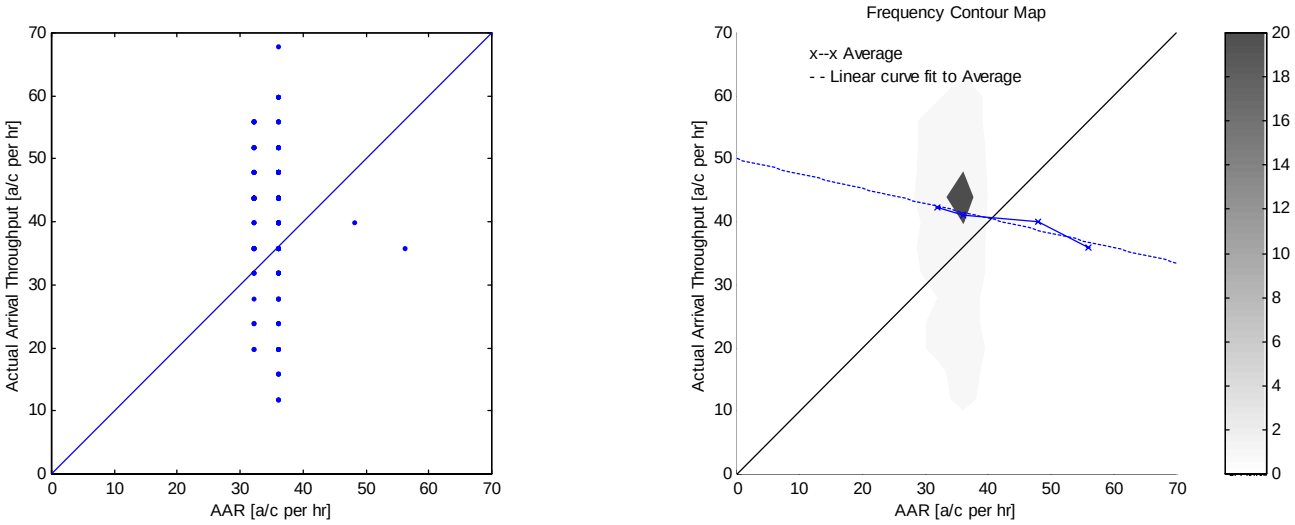

Figure C3. Airport utilization – raw data plot, and contour frequency plot.

Underutilization AAR	41 a/c per hr

C2.2. All JFK Configurations

The following plot shows all data points, with no filtering for high demand only:

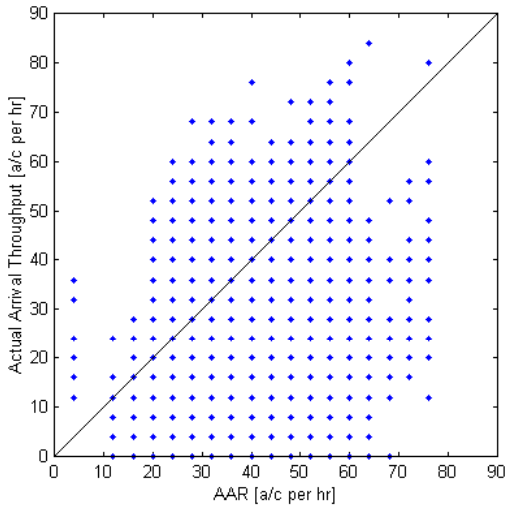

Figure C4. Airport utilization – raw data plot.

C3. LaGuardia Airport LGA

C3.1. Configuration 3 – 22 | 13

In the following chart actual arrival throughput is plotted against specified AAR in order to identify the degree to which the available capacity is being utilized. Only data with demand high enough to utilize the airport's maximum capacity is plotted. The threshold according to which demand is filtered is 52 aircraft per hour.

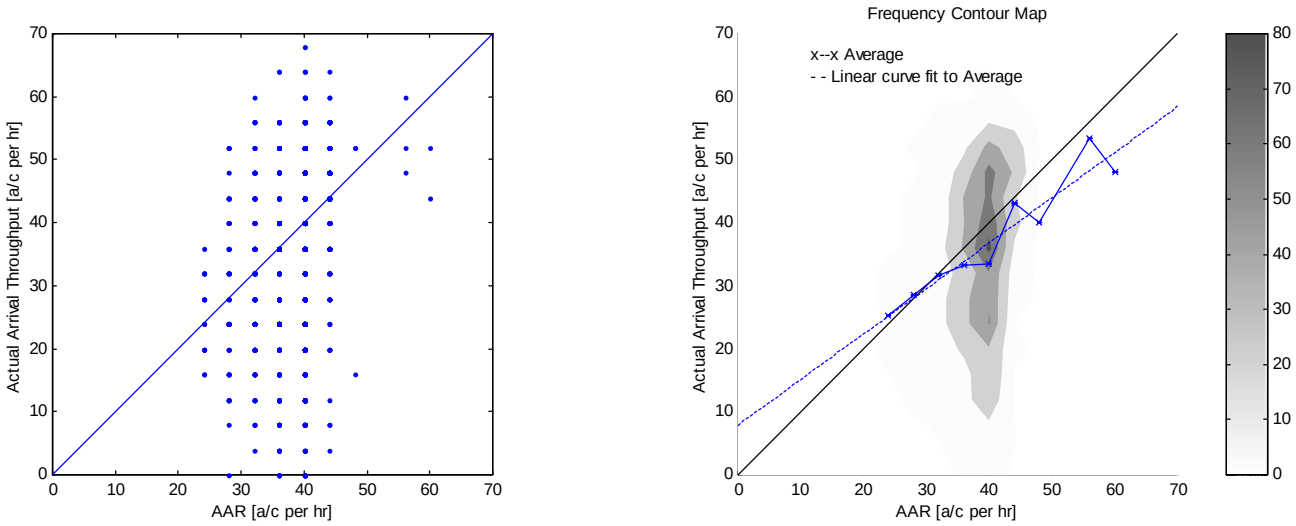

Figure C5. Airport utilization – raw data plot, and contour frequency plot.

Underutilization AAR	29 a/c per hr

C3.2. All LGA Configurations

The following plot shows all data points, with no filtering for high demand only:

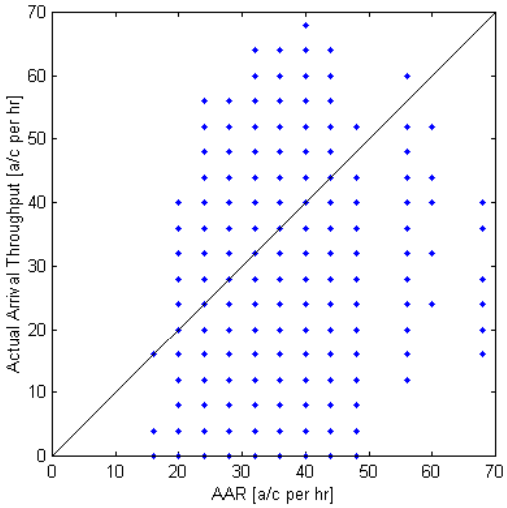

Figure C6. Airport utilization – raw data plot.

C4. Philadelphia International Airport PHL

C4.1. Configuration 4 – 26, 27R, 35 | 27L, 35

In the following chart actual arrival throughput is plotted against specified AAR in order to identify the degree to which the available capacity is being utilized. Only data with demand high enough to utilize the airport's maximum capacity is plotted. The threshold according to which demand is filtered is 68 aircraft per hour.

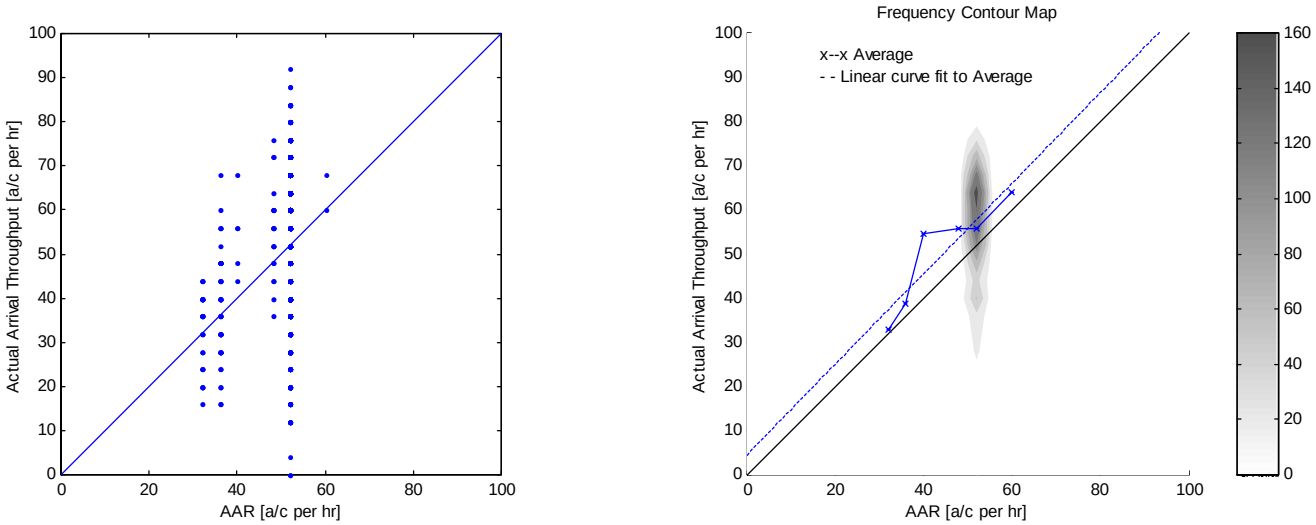

Figure C7. Airport utilization – raw data plot, and contour frequency plot.

Underutilization AAR	None

C4.2. All PHL Configurations

The following plot shows all data points, with no filtering for high demand only:

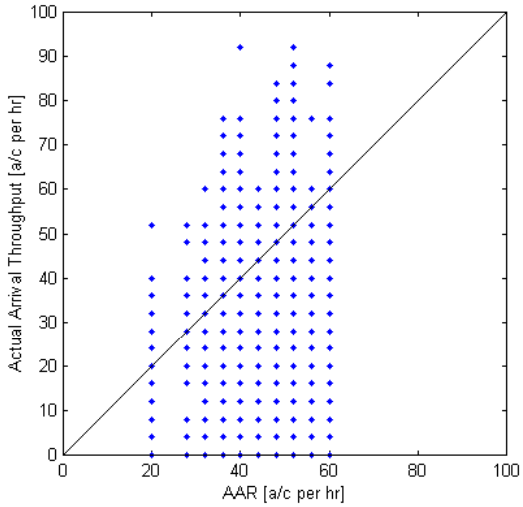

Figure C8. Airport utilization – raw data plot.

C5. Teterboro Airport TEB

Configurations were not detailed in the data on TEB. Consequently all data shall be analyzed as if it were a single configuration.

In the following chart actual arrival throughput is plotted against specified AAR in order to identify the degree to which the available capacity is being utilized. Only data with demand high enough to utilize the airport's maximum capacity is plotted. The threshold according to which demand is filtered is 28 aircraft per hour.

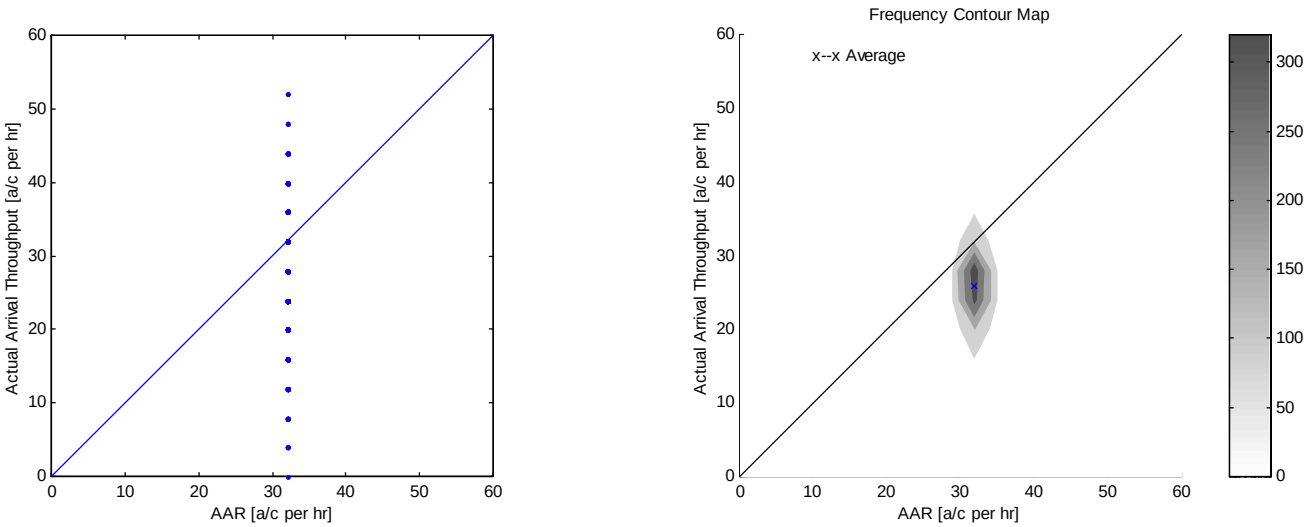

Figure C9. Airport utilization – raw data plot, and contour frequency plot.

Underutilization AAR	Only one AAR specified

C6. New York TRACON N90

The data from EWR, JFK, LGA and TEB is combined for those days for which all data sets have data. All this data is plotted as different combinations of configurations are not identified.

In the following chart actual arrival throughput is plotted against specified AAR in order to identify the degree to which the available capacity is being utilized. Only data with demand high enough to utilize each airport's maximum capacity (according to the configurations analyzed above) is plotted. The threshold according to which demand is filtered is 148 aircraft per hour.

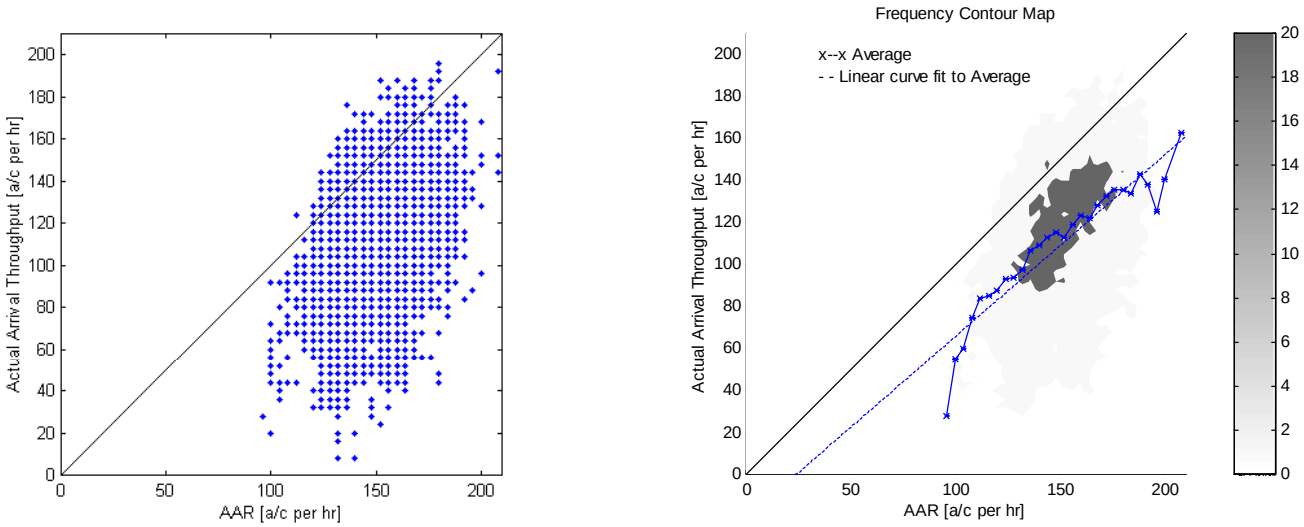

Figure C10. Airport utilization – raw data plot, and contour frequency plot.

Underutilization AAR	All

The following plot shows all data points, with no filtering for high demand only:

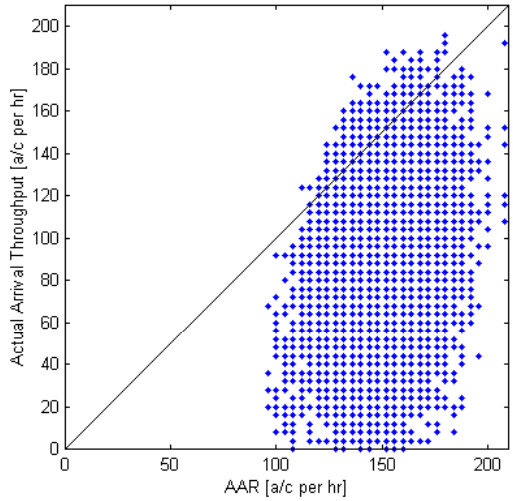

Figure C11. Airport utilization – raw data plot.

C7. Observations

A number of observations were made from the plots of arrival throughput versus AAR. These were as follows:

1. The majority of data points fall close to the line of 100% utilization in almost all cases (except TEB), when low demand is filtered out.
2. As AAR increases average utilization decreases from being greater than 1 for low AAR to being less than 1 for high AAR, for all airports.
3. For N90 alone is average utilization less than 1 for all AARs. This is in contrast to utilization at the airports at low AAR, which are greater than 1. This suggests that when the AAR at one airport is low, and the throughput high, the throughput at other airports within N90 is low. This suggests that there is an interaction between the airports within N90, and that the N90 airspace is a constraint.
4. All airports within N90 are over-utilized (throughput higher than AAR) by a significant degree in some periods. N90 however, although over-utilized during some periods, is over-utilized by only a small degree. This suggests that the AARs specified are more representative of the capacity of N90, and not of the individual airports. This suggests that N90 is the primary flow constraint, and not the airports, as suggested through interviews with TMCs at N90 in November, 2002. If the capacity of N90 were to be increased, it is thus suspected that the AARs specified at each airport could be increased. This increase would vary from airport to airport.

Appendix D: ASPM Analysis of Capacity Envelopes

Following is a capacity analysis of the four primary airports in New York (EWR, JFK, LGA and TEB), and of Philadelphia airport (PHL). As in Appendix B and C, the analysis is performed using the Aviation System Performance Metrics (ASPM) database, obtained through NASA and the FAA (http://www.apo.data.faa.gov/faamatsall.HTM). The results of the analysis are separated by airport. In order to accurately identify the capacity envelopes of the airports the data analyzed must include the periods of highest demand. Consequently data from January to August of 2001 was chosen for processing. The summer months are the period of highest demand, so data from the summer of 2000 was also included to enhance the data from the summer of 2001.

D1. Newark International Airport EWR

Capacity envelopes are plotted for a single airport configuration below. The configuration plotted is one of the commonly operated configurations at EWR.

D1.1. Configuration 3 – 22L | 22R

Capacity Envelopes

Arrival throughput is plotted against departure throughput in the figures below in order to identify the capacity envelopes for the airport.

In the charts quarter hourly arrival throughput is plotted against departure throughput. All data points are plotted. In order to represent rates by the hour the values for throughput and demand are multiplied by four.

The frequency of the data presented in the first chart is plotted alongside it, where the gray scale represents frequency. The frequency is again calculated in 4 aircraft per hour by 4 aircraft per hours bins.

Finally, below these charts, capacity envelopes are presented for the 100^{th} to 80^{th} percentile, in 4% increments. These percentiles are calculated for arrival throughput in 4 aircraft per hour bins of departure throughput. This allows departure throughput to be specified, and a corresponding arrival throughput to be identified from the envelope.

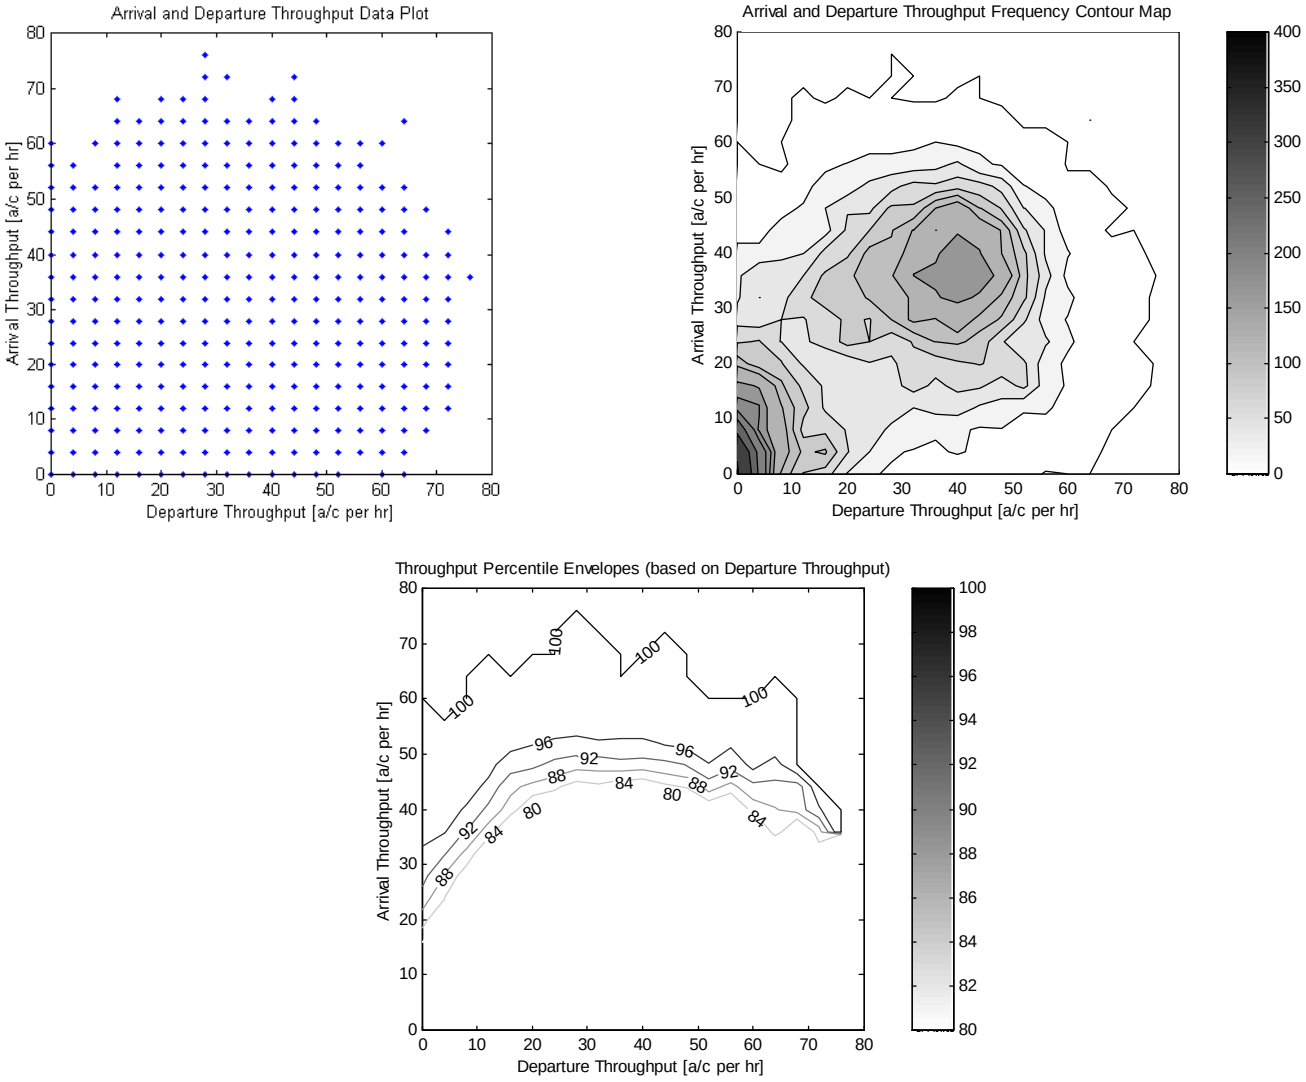

Figure D1. Airport capacity envelopes – raw data plot, contour frequency plot, and percentile envelopes.

D2. John F Kennedy International Airport JFK

Capacity envelopes are plotted for a single airport configuration below. The configuration plotted is one of the commonly operated configurations at JFK.

D2.1 Configuration 2 – 13L | 13R

Capacity Envelopes

Arrival throughput is plotted against departure throughput in the figures below in order to identify the capacity envelopes for the airport.

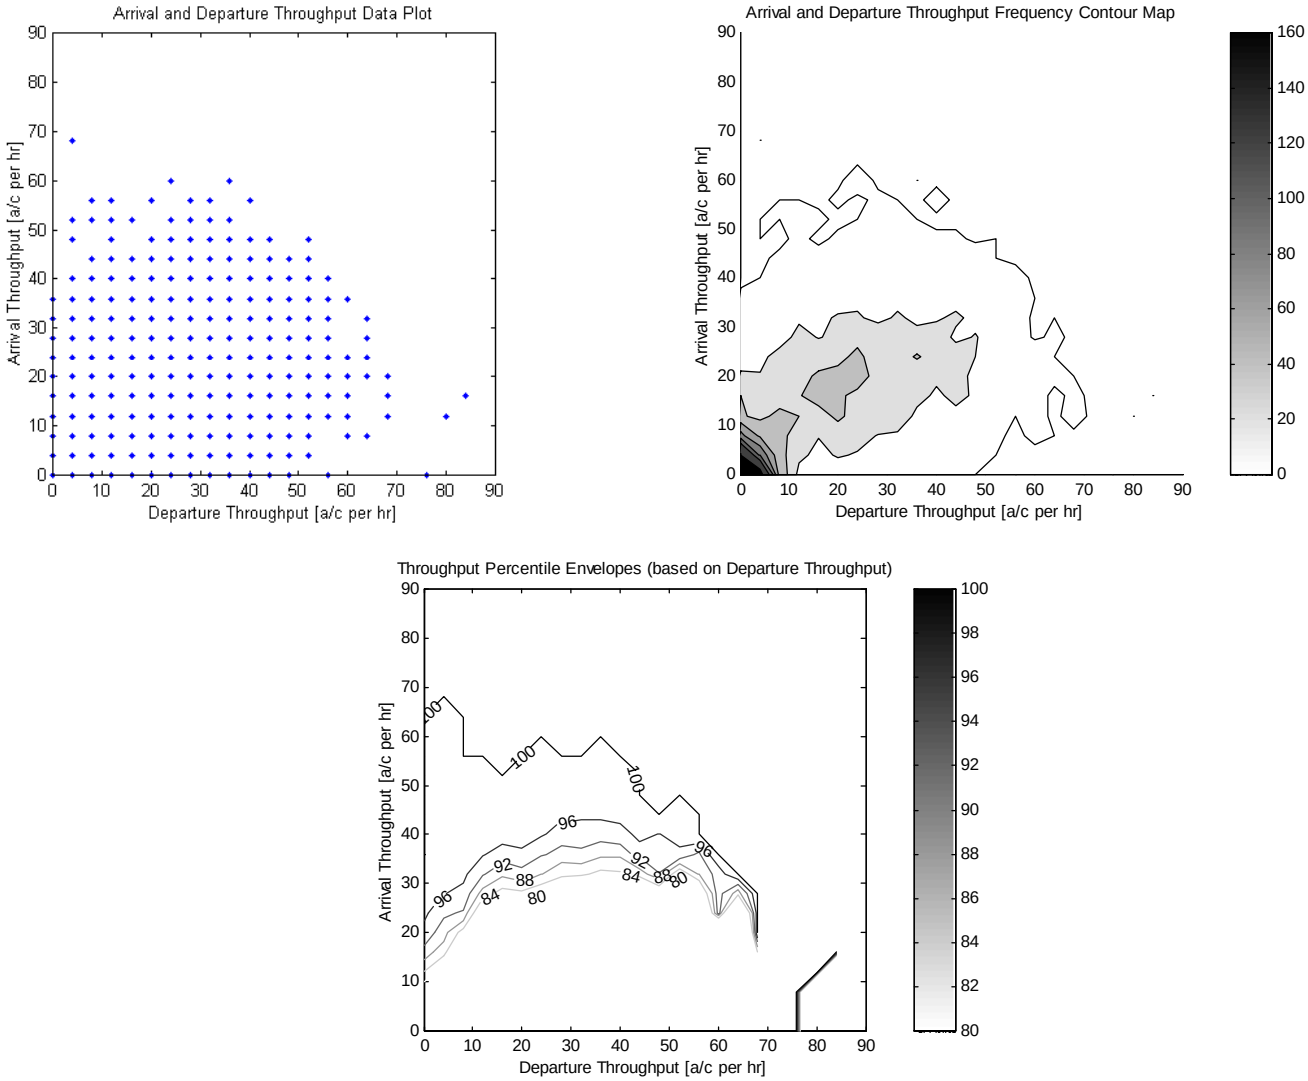

Figure D2. Airport capacity envelopes – raw data plot, contour frequency plot, and percentile envelopes.

D3. LaGuardia Airport LGA

Capacity envelopes are plotted for a single airport configuration below. The configuration plotted is one of the commonly operated configurations at LGA.

D3.1. Configuration 1 – 13 | 13

Capacity Envelopes

Arrival throughput is plotted against departure throughput in the figures below in order to identify the capacity envelopes for the airport.

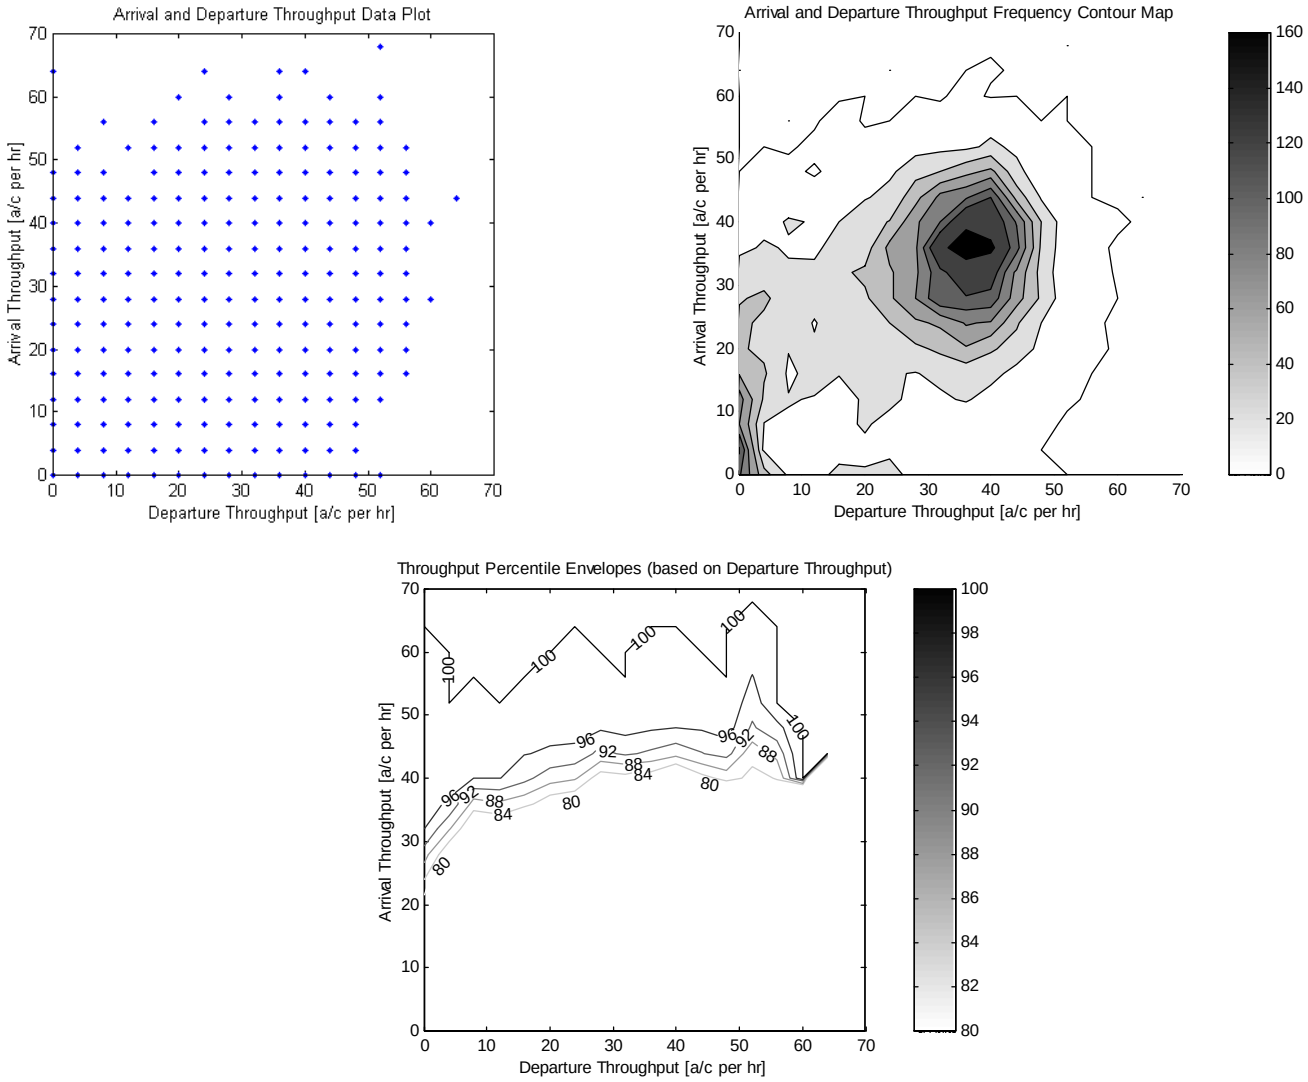

Figure D3. Airport capacity envelopes – raw data plot, contour frequency plot, and percentile envelopes.

D4. Philadelphia International Airport PHL

Capacity envelopes are plotted for a single airport configuration below. The configuration plotted is one of the commonly operated configurations at PHL.

D4.1. Configuration 4 – 26, 27R, 35 | 27L, 35

Capacity Envelopes

Arrival throughput is plotted against departure throughput in the figures below in order to identify the capacity envelopes for the airport.

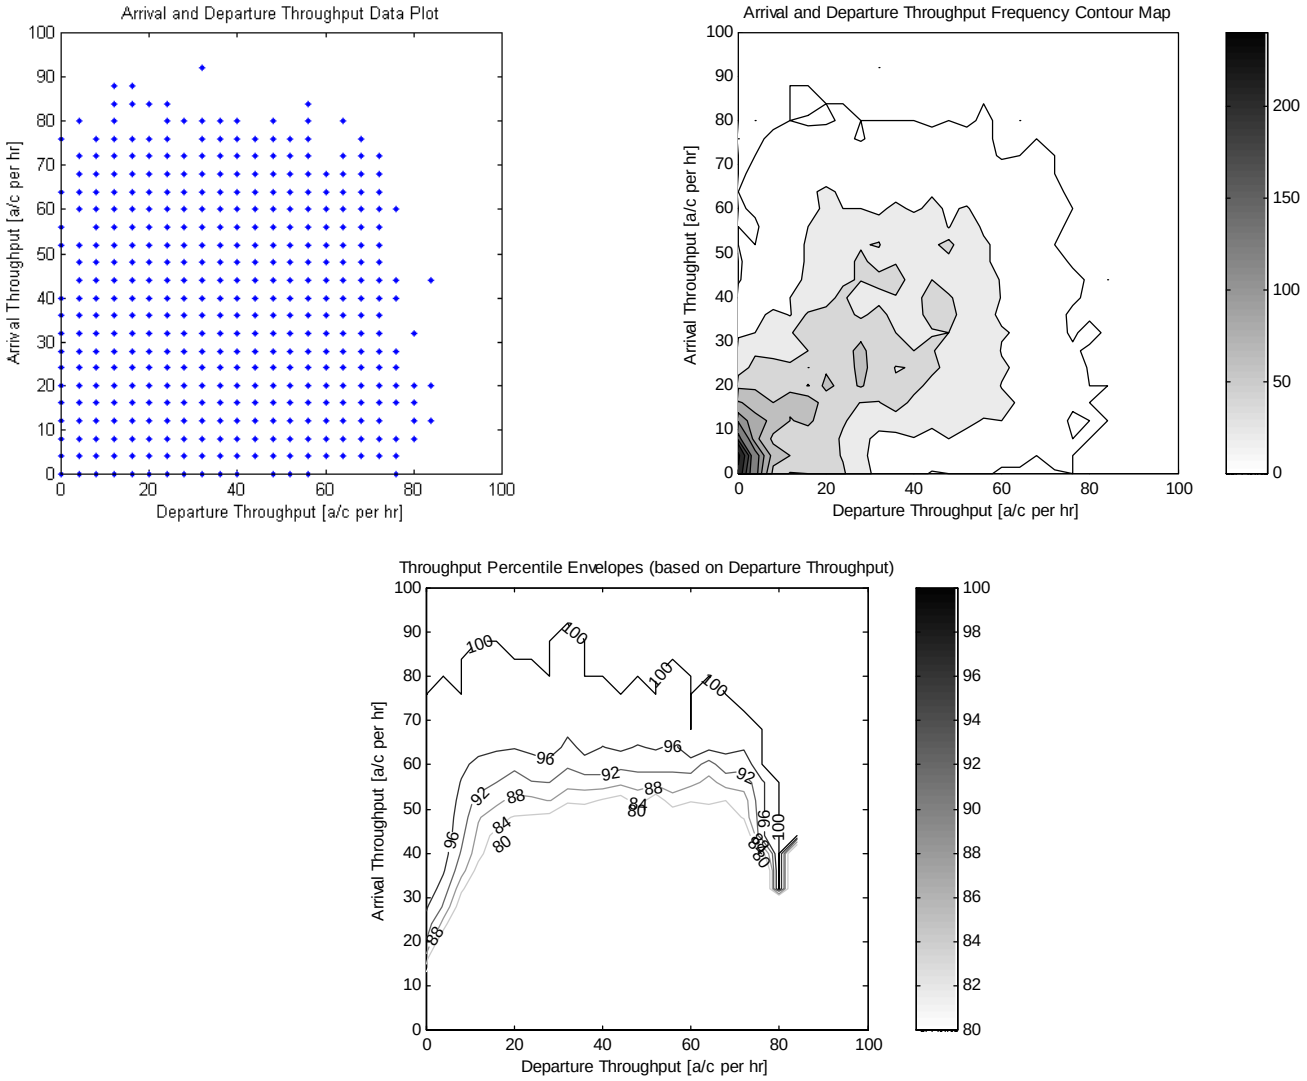

Figure D4. Airport capacity envelopes – raw data plot, contour frequency plot, and percentile envelopes.

D5. Teterboro Airport TEB

Configurations were not detailed in the data on TEB. Consequently all data shall be analyzed as if it were a single configuration.

Capacity Envelopes

Arrival throughput is plotted against departure throughput in the figures below in order to identify the capacity envelopes for the airport.

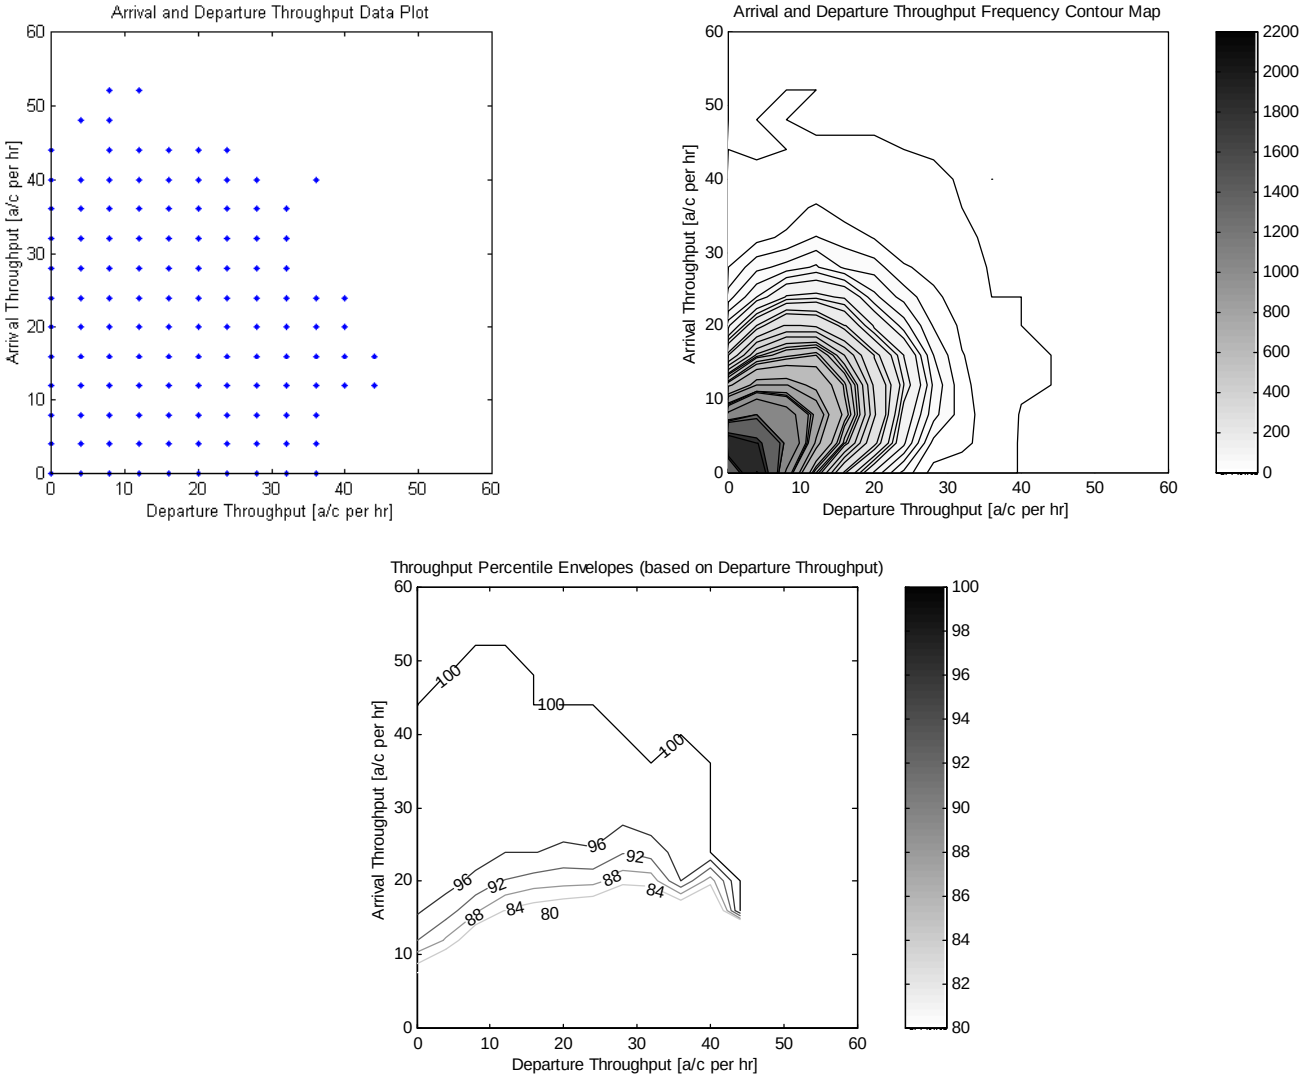

Figure D5. Airport capacity envelopes – raw data plot, contour frequency plot, and percentile envelopes.

D6. New York TRACON N90

The data from EWR, JFK, LGA and TEB is combined for those days for which all data sets have data. All this data is plotted because different combinations of configurations are not identified.

Capacity Envelopes

Arrival throughput is plotted against departure throughput in the figures below in order to identify the capacity envelopes for the airport.

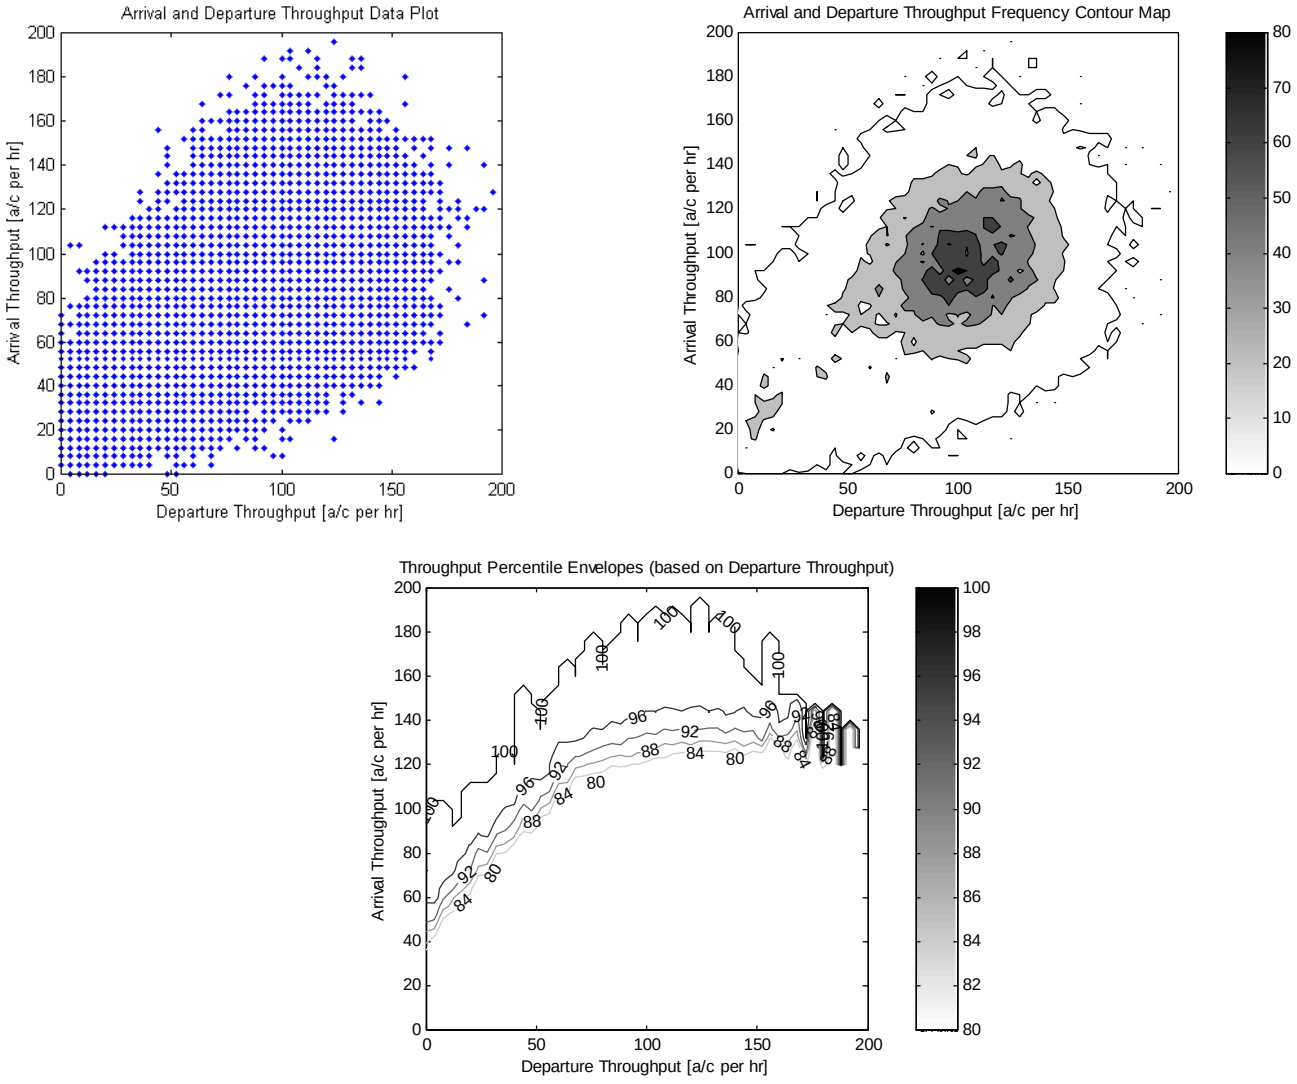

Figure D6. Airport capacity envelopes – raw data plot, contour frequency plot, and percentile envelope.